TÉCNICO LISBOA

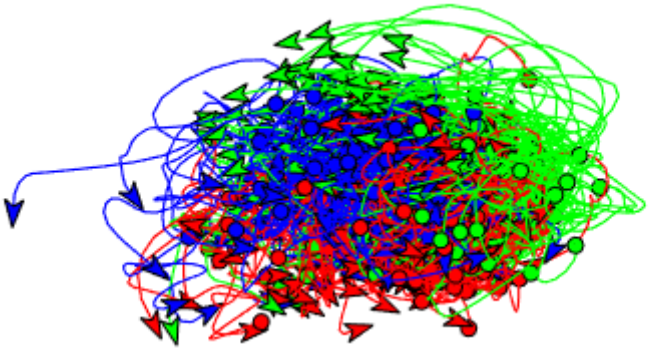

A dynamical approach to modelling visuomotor activity

Comparison between dorsal and ventral premotor cortices

Pedro Miguel Correia Pinto Sabido

Thesis to obtain the Master of Science degree in

Biomedical Engineering

Supervisors: Prof. Peter Janssen
Prof. João Sanches

October 2019

Very active, hard worker and independent student.
The thesis is very well written (and spoken), illustrated and supported by a solid and complete list of references

Publications?

Rotação: pouco motivado

Focar este trabalho em BMI não será limitativo?

I declare that this document is an original work of my own authorship and that it fulfils all the requirements of the Code of Conduct and Good Practices of the Universidade de Lisboa.

Preface

The work presented in this thesis was performed at the Laboratory for Neuro and Psychophysiology of the department of Neurosciences of the Katholieke Universiteit Leuven (Leuven, Belgium), during the period February-June 2019, under the supervision of Prof. Peter Janssen, and within the frame of the Erasmus programme. The thesis was co-supervised at Instituto Superior Técnico by Prof. João Sanches.

Acknowledgments

I would like to thank everyone that somehow contributed to my emotional and physical well-being and my academical success during this journey, that university is!

Primeiro que tudo, a minha família! Quero agradecer à minha mãe, por fazeres tudo o que consegues para me ver bem, e ao meu pai, por apoiares em todos os sonhos que tenho. Aos meus pais e avós quero dar um abraço individual, pela amizade, encorajamento, e amor desde sempre. Quero agradecer ao meu irmão pelo orgulho que ele nunca escondeu sentir e por sempre me querer proteger. Quero também agradecer à minha sobrinha e à Rosa pelo carinho e pelo apoio que me têm dado. A todas vocês: este trabalho é um bocadinho vosso (mas não muito, que foi a mim que isto tirou o sono!!!), por me terem sempre dado um espaço em que sou amado, apoiado e valorizado, e por me terem ensinado, muitas vezes sem saberem, como ser melhor.

I am so very thankful to my promotor, Peter Janseen, for the opportunity to join his lab and for his orientation, together with Elsie Premereur, both sharing their time and knowledge to help me develop the best work possible. I cannot mention my experience in Leuven without thanking Lara Merken, Sofie De Schijver and Kentaro Ogawa, who were super supportive friends during my stay in Leuven. A huge thank you to Wouter Depuydt, who was always, always available to help me with my IT problems.

Quero também agradecer ao meu orientador João Sanches pela disponibilidade que sempre teve e pelos esclarecimentos e orientação que contribuiram para o resultado final.

Não quero deixar de referir todos os amigos que fiz ao longo do curso, e que muitas vezes foram grande parte da motivação para sair da cama, para dar sempre o meu melhor em cada trabalho, que se riram comigo (e de mim!), e que me valorizaram como colega e como amigo: Helena Mendes, Marta Alves, Joana Real, Mariana Santos, João Silva, Bárbara Costa, Cláudia Dinis, e todos os meus colegas que fizeram parte do meu dia a dia. Um agradecimento especial à Carolin(d)a Barata, por ser parte da minha família austríaca de erasmus.

Não menos importantes são as pessoas com quem vivi na residência, e com quem partilhei a cozinha. Obrigado por serem umas porcas, mas por criarem um ambiente super cool, uma

verdadeira segunda casa.

Ich möchte auch mich bei Didem Yorgancıgil, Johannes Stöllinger und Kurt Hlavaty bedanken. Ihr habt mir das Gefühl gegeben, zu Hause zu sein, wann ich in Österreich war. Und es freut mich sehr, dass wir immer noch so oft sprechen, und enge Freunde sind.

Quero também agradecer ao João Gonçalves, que pode não ter o melhor feitio, pode ignorar metade do que eu digo, pode até testar a minha paciência constantemente, mas está lá sempre. Um obrigado do tamanho de Hong Kong ao Diogo Gonçalves por me ter ajudado diversas vezes nos meu infoproblemas e por querer sempre ouvir-me reclamar sobre a tese. Je dois aussi remercier ma amie et compagnie de voyages, 马芳达, 我很喜欢见到你， 你是我的女版。

To each and every one of you, here is my acknowledgement in all of your languages. – Obrigado. Thank you. Danke schön. Merci. 谢谢。Teşekkür ederim. Děkuji.

Abstract

Currently, intracortical BMIs rely exclusively on activity in the primary motor cortex. In many clinical conditions, however, this area is damaged or severely degenerated, which limits the restoration of movement to these patients. The premotor and parietal cortices, on the other hand, are less often affected and, as they coordinate the planning of motor commands through visuomotor integration, they represent a natural alternative for BMI control. Recent research also suggests that motor information can be extracted from the rotational structure underlying the temporal sequence of neuronal states in motor areas. The present work compares the adequacy of representing the neuronal activity in the dorsal (PMd) and ventral (PMv) premotor cortices during a reach-to-grasp task with such a dynamical model. The results show very robust dynamical rotations in both areas in all epochs of the task, despite the low response in PMv before movement onset. Furthermore, this project also seeks to investigate the usability and flexibility of a visuomotor BMI that uses the dynamical rotations of premotor neuronal populations to control movement direction and onset. The decoding accuracy of both these parameters is very much above the chance level, with even better performance than the decoding without the application of this model. This work suggests that a dynamical representation of premotor activity may yield satisfactory results for the control of a motor BMI in a discrete space.

Keywords

Visuomotor BMI; state-space modelling; neuronal dynamics; rotational dynamics in PMd and PMv; decoding of motor parameters.

Resumo

Atualmente, as interfaces cérebro-máquina (ICM) intracorticais recorrem exclusivamente à atividade no córtex motor primário. No entanto, quadros clínicos em que esta área cortical esteja danificada ou severamente degenerada comprometem o restabelecimento da atividade motora aos doentes. Os córtexes premotor e pariental, por outro lado, são menos frequentemente afetados, e, uma vez que estão envolvidas no planeamento da atividade motora através de integração visuomotora, representam alternativas naturais para o controlo de ICMs motoras. Estudos recentes sugerem ainda que a informação motora pode ser extraída da estrutura rotacional que carateriza a sequência temporal de estados neuronais em regiões motoras. Este trabalho pretende comparar a qualidade da representação da atividade neuronal nos córtexes premotores dorsal (PMd) e ventral (PMv) durante uma tarefa de alcaçar-para-agarrar através de um modelo dinâmico. Os resultados revelam uma estrutura rotacional robusta em ambas as áreas em todas as épocas do movimento, apesar da baixa responsividade em PMv antes do início do movimento. Adicionalmente, este projeto investiga a usabilidade e flexibilidade de uma ICM que use rotações dinâmicas de populações de neurónios premotores para controlar a direção e o início do movimento. A precisão de decodificação de ambos os parâmetros com este modelo dinâmico é bastante superior ao nível de chance, e está acima dos resultados obtidos para a decodificação feita sem a aplicação do modelo. Este trabalho sugere que uma representação dinâmica da atividade premotora pode produzir resultados satisfatórios para controlar uma ICM motora num espaço discreto.

Palavras-chave

Interface cérebro-máquina com input visuomotor; modelação no espaço de estados; dinâmica neuronal; rotações dinâmicas em PMd e PMv; decodificação de parâmetros motores.

Contents

1 Introduction **1**

1.1 Motivation . . . 1

1.2 Objectives . . . 2

1.3 Project Outline . . . 3

2 Scientific Background **5**

2.1 Visuomotor encoding in the brain . . . 5

2.1.1 Ventral premotor cortex, F5 (PMV) . . . 8

2.1.2 Dorsal premotor cortex, F2 (PMd) . . . 8

2.2 Brain-machine Interfaces (BMIs) . . . 9

2.2.1 Evolution of BMI research . . . 9

2.2.2 Electrophysiological recordings . . . 11

2.2.3 Information encoding by neurons and neuronal ensembles . . . 12

2.2.4 Modelling neuronal activity . . . 13

2.2.5 Decoding of neuronal signals . . . 15

3 Methods **17**

3.1 Neuronal data acquisition . . . 17

3.1.1 Electrophysiological recordings . . . 17

3.1.2 Experimental paradigm . . . 19

3.1.3 Data Processing . . . 19

3.2 jPCA . . . 20

3.2.1 Mathematical description . . . 21

3.2.1.A Solving for the generic linear dynamical system . . . 21

3.2.1.B Solving for the constrained linear dynamical system . . . 22

3.2.2 Variance captured by a jPCA plane . . . 24

3.2.3 Goodness of fit . . . 24

3.3 Support vector machine . . . 25

3.3.1 Decoding movement direction . . . 25

3.3.2 Decoding movement onset . . . 26
3.4 Material Specifications . . . 26
4 Results **27**
4.1 Neuronal activity in PMd and PMv . . . 27
4.1.1 From spikes to spike rates . . . 27
4.1.2 Common response profiles . . . 28
4.2 Rotational structure in PMd and PMv dynamics . . . 31
4.2.1 State-space modelling with jPCA . . . 32
4.2.2 Goodness of fit of the jPCA model . . . 40
4.2.3 Control analyses . . . 41
4.2.3.A Disrupting the dynamical structure . . . 42
4.2.3.B Altering the data covariance . . . 45
4.3 Unraveling the neuronal code in PMd . . . 46
4.3.1 Patterns in neuronal activity . . . 46
4.3.2 Decoding movement direction . . . 48
4.3.2.A Input to Support Vector Machine (SVM) . . . 48
4.3.2.B Decoding accuracy . . . 51
4.3.3 Decoding movement onset . . . 53
4.3.3.A Input to Support Vector Machine (SVM) . . . 53
4.3.3.B Decoding accuracy . . . 54
5 Discussion **57**
5.1 Neuronal activity in PMd and PMv . . . 57
5.2 Rotational structure in PMd and PMv dynamics . . . 58
5.3 Unravelling the neuronal code in PMd . . . 60
6 Conclusion **65**
6.1 Functional properties of the premotor cortex . . . 65
6.2 Methodological considerations . . . 67
6.3 Future perspectives for real-time BMI control . . . 68
References **69**
A Geometrical interpretation of jPCA **77**
B State-space representation of disrupted data **82**

List of Figures

2.1 Cortical areas of the brain involved in the planning and execution of visuomotor tasks . . . 6
2.2 Simplified description of interactions between cortical structures during a visuomotor task. . . . 7
2.3 Components of a brain-machine interface for motor control . . . 10
2.4 Utah Array . . . 12
2.5 Oscillation of neural firing rates during three movement types. . . . 15

3.1 Location of the chronically implanted Utah arrays in the second experiment. . . . 18
3.2 Schema of the delayed reach-to-grasp task. . . . 20

4.1 Temporal raster plot of the spikes detected in channel 16 of the array implanted in PMd in monkey J3-PMd. . . . 28
4.2 Unfiltered and Gaussian-filtered PSTH of the multi-unit activity detected in channel 16 of the array implanted in PMd of monkey J3 during a reach-to-grasp task. . . . 29
4.3 Across-channel average PSTH of the multi-unit activity recorded in all 64 channels in PMd and PMv of monkey J3. . . . 29
4.4 Filtered PSTH of the multi-unit activity detected in PMd of monkey J3 during a reach-to-grasp task. . . . 30
4.5 Filtered PSTH of the multi-unit activity detected in PMv of monkey J1 during a reach-to-grasp task. . . . 31
4.6 Projection onto state-space of the multi-unit activity of a PMd population in monkey J3 during 5 intervals of interest of a reach-to-grasp task. . . . 34
4.7 Projection onto state-space of the multi-unit activity of a PMv population in monkey J3 during 5 intervals of interest of a reach-to-grasp task. . . . 37
4.8 Comparison between PMd and PMv of the variance captured by the 6 jPCA projections found within the 6-dimensional state-space obtained via PCA. . . . 40

4.9 Comparison between PMd and PMv of the coefficient of determination of the constrained system, $R^2_{M_{skew}}$, and of the coefficient of determination ratio between the constrained and the unconstrained systems $R^2_{M_{skew}}/R^2_M$. This analysis was performed for 5 different intervals of interest. 41
4.10 Histograms of the angle between the neural state, X, and its derivative, $\dot{X}$. 43
4.11 Comparison between PMd and PMv of the effect of the shuffled controls on the coefficient of determination of the constrained system, $R^2_{M_{skew}}$, and on the ratio between the coefficient of determination of the constrained and the unconstrained system $R^2_{M_{skew}}/R^2_M$. This analysis was performed for datasets J1-6, during 200 ms after object onset. 44
4.12 Comparison of the statistical significance of the coefficients of determination of the constrained models (and their ratio with the unconstrained models) between PMd and PMv, in 1000 bootstraps. 45
4.13 Neuronal activity recorded for all trials in a given dataset for four different channels. 47
4.14 Projection onto state-space of all the trials recorded in monkey K8, referring to the multi-unit activity of a PMd population, for three intervals. 49
4.15 Examples of the input observations used to decode and control the decoding of movement direction. 50
4.16 Accuracy of decoding movement direction with jPCA projections versus the entire neuronal ensemble. 51
4.17 Time necessary to train the SVM for decoding movement direction with jPCA projections versus the entire neuronal ensemble. 52
4.18 Examples of the input observations used to decode and control the decoding of movement onset. 54
4.19 Accuracy of decoding movement onset with jPCA projections versus the entire neuronal ensemble. 55
4.20 Time necessary to train the SVM for decoding movement direction with jPCA projections verus the entire neuronal ensemble. 55

A.1 3-Dimensional artificial data used in jPCA explanation 77
A.2 Result of performing PCA on the artificial data 78
A.3 Projection of the artificial data onto the two eigenvectors of M_skew and M_sym with the biggest eigenvalues. 80
A.4 Projection of data onto a plane defined by the eigenvectors of a symmetric summary matrix. 80

A.5 Projection of data onto a jPCA plane defined by the eigenvectors of a skew-symmetric summary matrix. 81

B.1 Effect of the shuffle controls on the PSTH of the multi-unit activity recorded in channel 8 of the array implanted in PMd of monkey J3. 83
B.2 Effect of the shuffle controls on the PSTH of the multi-unit activity recorded in channel 8 of the array implanted in PMv of monkey J3. 84
B.3 Effect of the shuffle controls on the jPCA projections of the multi-unit activity of a PMd population in monkey J3 during the 200 ms after object onset. 85
B.4 Effect of the shuffle controls on the jPCA projections of the multi-unit activity of a PMv population in monkey J3 during the 200 ms after object onset. 87

List of Tables

3.1 Channels of the Utah arrays recorded in PMd and PMv of monkey J in the second experiment. 18

4.1 Intervals used in the comparative analysis of the dynamical structure. The column ≪Alignment≫ defines the event to which the starting and ending points are relative to. 32
4.2 Intervals used in the decoding of movement direction. The column ≪Alignment≫ defines the event to which the starting and ending points are relative to. 48
4.3 Intervals used in the decoding of movement onset. The column ≪Alignment≫ defines the event to which the starting and ending points are relative to. 53

List of Acronyms

BMI	Brain-machine Interface
PCA	Principal Componenent Analysis
PC	Principal Component
jPC	jPCA projection
IPS	intraparietal sulcus
ITC	inferotemporal cortex
AIP	anterior intraparietal area
PIP	posterior intraparietal area
PFC	prefrontal cortex
VIP	ventral intraparietal area
M1	primary motor cortex
SMA	supplementary motor area
pre-SMA	pre-supplementary motor area
PMd	dorsal premotor area
pre-PMd	pre-dorsal premotor cortex
PMv	ventral premotor area
LFP	local field potential
MUA	multi-unit activity
SUA	single-unit activity

PSTH Peristimulus Time Histogram

CD coefficient of determination

RGR rotational goodness-of-fit ratio

SVM Support Vector Machine

LSTM Long Short-Term Memory

Chapter 1

Introduction

This chapter contextualises this project in the current needs of neuroscientific research. An explanation of the problem at hands is also included, as well as an overview of the structure of this thesis.

1.1 Motivation

Every conscious movement a person makes, whether playing a guitar or lifting a glass of water, begins in an area of the brain named motor cortex. The primary motor cortex (M1) is responsible for sending these motor commands along the spinal cord, which are passed onto motor neurons. In many clinical conditions, however, the limb may be amputated, the spinal cord might be injured, or the primary motor neurons may be damaged (e.g. after a stroke) or severely degenerated (e.g. in amyotrophic lateral sclerosis), compromising the patients' mobility and limiting clinical rehabilitation. (1,2)

Naturally, the restoration of movement to patients with some level of motor impairment can have a determinant impact in their biopsychosocial adjustment. With that in mind, research in Brain-machine Interfaces (BMIs) has become an emergent field of research. However, most major studies in humans have been conducted with M1 inputs, an area that can be affected in some of the aforementioned clinical conditions. (1,3)

An alternative to M1 motor inputs is visuomotor neural activity of the premotor areas. Both dorsal premotor area (PMd) and ventral premotor area (PMv) use information about object characteristics and spatial organisation to drive the planning of motor behaviour. (4–7) Therefore, potentially useful information for BMI control can be retrieved from these areas, mitigating the limitations caused by motor impairment.

1.2 Objectives

The classical view of the motor cortex holds that motor cortical neuronal activity represents movement or muscle parameters. However, recent research in humans and non-human primates has shown that the temporal structure of neuronal activity in M1 and PMd reveals orderly rotational structure that can be used to extract motor information. (2,8) Combating this lack of agreement and understanding the motor cortex is, thus, a major challenge, but an essential step toward designing more accurate and capable neuroprosthetic devices. (9)

The main goal of this work is to investigate the usability and flexibility of a visuomotor activity recorded in the dorsal (PMd, F2) and ventral (PMv, F5c) premotor cortices for motor control. This master's thesis aims to serve as foundation for the development of a novel visuomotor BMI for human applications, that does not rely on the actual motor execution command, but rather on neuronal activity underlying motor planning. However, before performing invasive procedures in the human brain, this study will be conducted in a monkey model, due to resemblance of their brain with ours.

The first question to answer is whether or not multi-unit activity (MUA) in PMv is well-described by a dynamical system, and if the temporal sequence of neuronal states in this motor area has a rotational structure. These results will be compared with those obtained for another neuronal population in PMd recorded simultaneously during a reach-to-grasp task. This analysis will confer a better understanding of how premotor areas relate to reaching and grasping tasks, while filling in the research gap in comparing the rotational dynamics that describe the MUA in both these areas.

These results will then be interpreted in the light of recent studies that connect the rotational structure of neuronal trajectories with movement and muscle parameters. This leads us to the second objective of this project, which is to evaluate the information that can be extracted from the rotational structure of the neuronal state, regarding movement direction and movement timing. It is worth noting that the purpose is not to determine the most accurate algorithm to retrieve information from rotational dynamics, but rather to investigate whether or not neuronal dynamics are a more efficient and accurate way of representing motor information than the classical spike rates.

The initial intent was to compare the decoding performance obtained from the dorsal and ventral premotor areas in the different epochs of the task (object onset, visual phase, and motor phase), and measure the stability of the decoding over a period of several weeks. However, technical complications determined an early end of the experiment, and the short duration of each recorded session precluded the use of these datasets for decoding. For these reasons, different datasets are going to be used for the two different parts of this thesis, one of which,

used for decoding, only included activity recorded in PMd .

1.3 Project Outline

This thesis starts with the introduction of relevant anatomical and functional aspects of the brain that may help understanding the neural connections that command prehension tasks, and movement in general. This chapter also presents the concept of BMIs, together with a historical perspective of neuroscientific research, which is particularly interesting in order to comprehend how the current understanding of neuronal coding has been shaped, and how machine learning bridges this knowledge with real life applications.

The next chapter outlines comprehensively the experimental procedure to acquire data from two premotor areas, PMd and PMv, of the study subject, a rhesus monkey, and the task performed. Also included is a detailed mathematical description of jPCA, a method developed to visualise the rotational patterns found in neuronal dynamics, as well as the statistical entities that are often used to evaluate it. Next, a description is included of how a Support Vector Machine (SVM) is implemented to perform two decoding tasks in discrete spaces.

The subsequent chapter describes all the results and tests performed. Firstly, a brief introduction of some of the responses found in the recorded datasets. These same datasets are then modelled with jPCA and various tests are conducted to conclude about the validity of the results. In an attempt to understand how rotations may encode information, two decoding tasks are held: one to decode task condition, or movement direction, and another to determine whether rotational dynamics could be used to signal motion initiation. All these results are subsequently discussed in light of the most recent literature.

The following and last chapter summarises the conclusions that can be drawn from this project and suggests areas of further research and improvements to this work that could potentially contribute to real life applications in neurorehabilitation.

Chapter 2

Scientific Background

For some, neuroprostheses and neurorehabilitation strategies might seem implausible in the near future, but they are here and their value to the biopsychosocial adjustment of patients is undeniable. This section overviews the basics of brain motor function and the evolution in brain-machine interface technology in the context of prehension tasks.

2.1 Visuomotor encoding in the brain

Humans and other primates are undeniably skilled at carrying out complex and fine movements, a facility that has contributed to the prosperity of our species. Our motor agility and dexterity is a reflection of our brain's capability to plan, coordinate and execute movement. (10) Therefore, the intimidating task of unravelling the physiological mechanisms that support the operation of the human brain has always generated substantial interest, and, thus, made the brain a target of various functional and anatomical research endeavours. (1)

In the past century, several anatomical studies to the cortex of primates led to functional subdivisions of the motor areas into primary motor (F1), premotor and supplementary motor areas, as represented in fig. 2.1. (4, 11) Since then, the premotor cortex is divided into a dorsal region, PMd, whose rostral (F7) and caudal (F2) regions are respectively termed pre-dorsal premotor cortex (pre-PMd) and PMd proper, and a ventral region, also parted into rostral (F5) and caudal (F4) parts. The supplementary motor region is also similarly subdivided, with a pre-supplementary motor area (pre-SMA) (F6) and a SMA (F3) proper region. (4)

In our daily activities, motor skills are often coordinated with, or in response to, visual input. Visuomotor skills describe the actions produced when visual and motor cortices work in concert. (14) It is, therefore, essential to understand how the representation of visual object

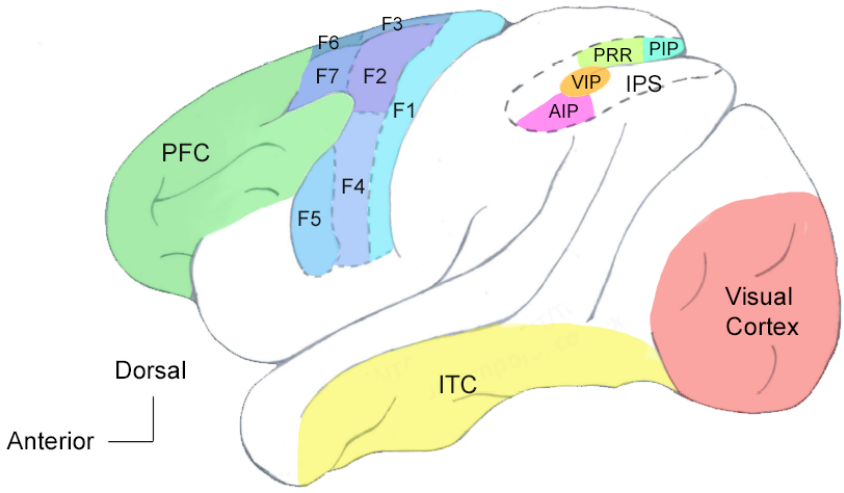

Figure 2.1: Schematic diagram of a monkey brain. Lateral view of the cortical areas of the left hemisphere involved in the planning and execution of visuomotor tasks. The traced area in the parietal cortex illustrates an unfolded view of the intraparietal sulcus. (Adapted from Janssen et al. (2018) (12) and Katsuki et al. (2013) (13))

properties transitions into intended motor acts - i.e, how the brain "sees" -, in order to figure out how the motor areas process reaching and grasping tasks.

Almost four decades ago, Ungerleider and Mishkin (1982) proposed the first model of the visual system, dividing it into a dorsal pathway, corresponding to the projections established from the visual cortex to several areas in the intraparietal sulcus (IPS) (in the posterior parietal cortex) and ventral stream, running from the visual cortex to the inferotemporal cortex (ITC). Several ensuing studies have corroborated this dichotomy, suggesting that the structures in the ventral stream are essential for processing object information for visual perception and object identification (replying to "what is the object?"), and the dorsal stream plays a role in the processing of spatial information and object characteristics to guide actions (giving an answer to "where is the object and how does it look?"). (12, 15, 16)

It comes with no surprise that, when grasping an object, having prior knowledge about its size, texture, predicted weight, purpose, usability, etc., is determinant to guide the grasping action and the pre-shaping of the hand. Inspired by studies in patients with occitotemporal or parietal lesions, Fagg and Arbib (1998) hypothesized a model of visually-guided grasping - the FARS (Fagg-Arbib-Rizzolatti-Sakata) model. (12)

This model establishes the anterior intraparietal area (AIP) as the first stage in the grasp programming process, where object-related information from the posterior intraparietal area (PIP) (dorsal stream) and object identification from the ITC (ventral stream) are integrated and a set of possible ways of grasping the objects - called affordances - is computed. This information is passed on to PMv (F5), where, given a set of constraints, only one grasp is selected. These

constraints include task information from pre-SMA (F6), working memory or behaviour inhibition (e.g. due to social and conventional rules) from the prefrontal cortex (PFC) - which relies on object recognition in ITC, and instruction stimuli from PMd (F2). In this last case, F2 is known to bias the selection of the grasp in tasks where the right grasp is conditional upon presentation of a stimulus. (15) PMv (F5) is then responsible for the execution and monitoring of the preshaping and grasping. The spatial information concerning the target object is represented in ventral intraparietal area (VIP) and transmitted to PMv (F4), where it produces an appropriate reach action. (15, 16)

A pictographic summary of all the aforementioned interactions between cortical structures during a reach-to-grasp task is represented in fig. 2.2. This schema should be analysed together with fig. 2.1, for an anatomical understanding.

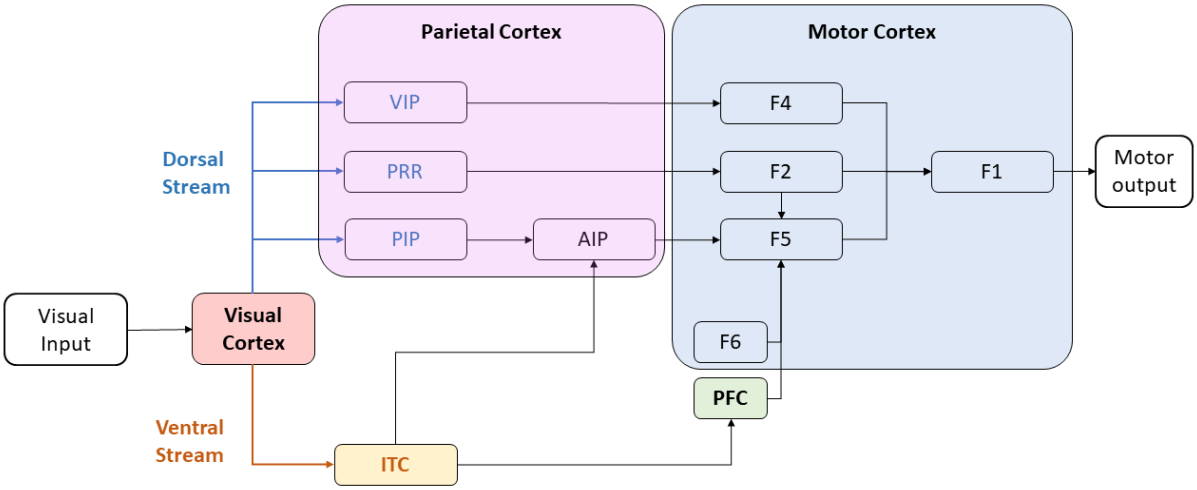

Figure 2.2: Simplified description of interactions between cortical structures during a visuomotor task. The schema is based on the FARS model and subsequent studies. The anatomical position of the cortical structures is identified in fig. 2.1. (Adapted from Fagg and Arbib (1998) fig. 2.1, Erhan Oztop (2002) (16) and Rizzolatti et al. (2003) (17).)

Although this dichotomy between reaching in PMd and grasping in PMv is strongly supported by and influential in recent research projects, this classical division is not absolute. In fact, Takashi et al. (2017) found grasping-tuned neurons in PMd, and neurons with a preference for reaching in PMv. (5)

This work focuses on the activity in PMd and PMv during a reach-to-grasp task. As such, the next two sections provide a better insight into these two premotor areas.

2.1.1 Ventral premotor cortex, F5 (PMV)

PMv (F5) has been subject to numerous electrophysiological studies that describe the determinant role of this area in visually-guided tasks. (7) Together with the AIP, F5 neurons transform the representation of the geometrical properties of the object into appropriate grasping commands, namely the shaping of the hand. (6)

Rizzolatti et al. (1988) described various groups of F5 neurons that discharge during specific hand movement, namely grasping, holding, tearing and manipulating, being the largest class related to grasping. (18) The author suggested also that F5 neurons have their own "motor vocabulary", showing selectivity to the type of hand grip required. i.e., precision grip, finger prehension, or whole-hand grasping. (19) Upholding these findings, Umilta et al. (2008) described F5 grasping neurons as goal-directed, rather than movement-directed. (20)

Given the cytoarchitectonics of F5, this area can be subdivided into three areas: F5 anterior (F5a), F5 posterior (F5p) and F5 convexity (F5c). (21) According to their functional properties, F5 grasping neurons can be classified into motor-dominant, visuomotor (or canonical), visual-dominant (almost exclusively in F5a) and "mirror" neurons (mainly in F5c). Motor-dominant neurons do not respond to the presentation of objects, but rather prepare a grasping action regardless of the lighting conditions. Visuomotor neurons discharge congruently with their motor specificity when objects are presented with a certain size, shape and orientation, even when there is no interaction with the object itself. (19) Visual-dominant neurons are active only during object fixation in the light, and are frequently selective for the depth structure of objects. (22) Lastly, the "mirror" neurons are active during both action execution as well as action observation. (23)

2.1.2 Dorsal premotor cortex, F2 (PMd)

While PMv seems to reflect visual target-location, crucial for object reaching and action understanding, PMd plays a more significant role in the planning of actions by retrieving, retaining and integrating visuospatial information about both target location and arm use. (7,24) In fact, Cisek and Kalaska (2005) and Hoshi and Tanji (2006) reported that, when primates face multiple potential reaching actions, the activity in PMd can reflect the planning of several directional signals, which are eliminated when a subsequent nonspatial cue identified the correct action. (7,25)

PMd (F2) has also been demonstrated to exhibit anticipatory activity for the forthcoming cue, with some neurons responding according to the motor set that must be prepared given a stimulus. (7,15) Churchland et al (2006) even reports that PMd activity can predict the timing

when an action takes place. (26)

F2 neurons have also been classified according to the timing of their activity in the context of motor tasks: some F2 neurons discharge in response to a visual stimulus, other neurons have movement-related activity, with a strong selectivity for the direction of the movement, and a third group of neurons fire during a delay period in anticipation for the go cue. (27)

It has been demonstrated that PMd has a more complete anatomical representation of the upper limb than PMv. Moreover, research supports the existence of a stronger representation of movement kinematics in PMd compared with PMv. (5)

2.2 Brain-machine Interfaces (BMIs)

For centuries, people have tried to understand the principles of behaviour of the dynamic circuits of the brain with the aim to apply them in the development of new therapies and the restoration of mobility and sensation to disabled patients. Eventually, the term BMI was coined to describe a direct communication pathway between the brain of an organism and an artificial actuator, like a virtual or robotic arm. BMIs allow subjects to utilise their own volitional electrical brain activity to control the motion of said actuator, while receiving continuous visual feedback. Currently, BMIs can also deliver sensory feedback from external actuators back to the brain. (1)

The design of a functional BMI requires three fundamental components: electrophysiological recordings, decoding algorithms and robotics. Neuronal signals are used as input data for the decoding software that interprets the information - e.g. predicts an intended action - and redirects it to the computer to acts upon the command – e.g. executing a movement. fig. 2.3 illustrates the deconstruction of a BMI for motor control. (1)

2.2.1 Evolution of BMI research

The history of BMIs is tightly connected to the work put into developing new neurophysiological technologies to record the electrical activity of the brain. (1)

Ramón y Cajal (1997) marked the emergence of modern neuroscience with their neuron doctrine. This theory established individual neurons as the functional unit of the brain, responsible for processing and transmission of electrophysiological signals. (29) In 1949, Donald Hebb presented the idea that the brain is wired as a neuronal network, and information is not encoded on the action of specific cells, but rather on the activity of neuronal populations. (1,30) However, multi-channel recordings did not become a reality until the 1990s, when the modern concept of BMI was introduced. The first experiments were carried out in rodents, and only in

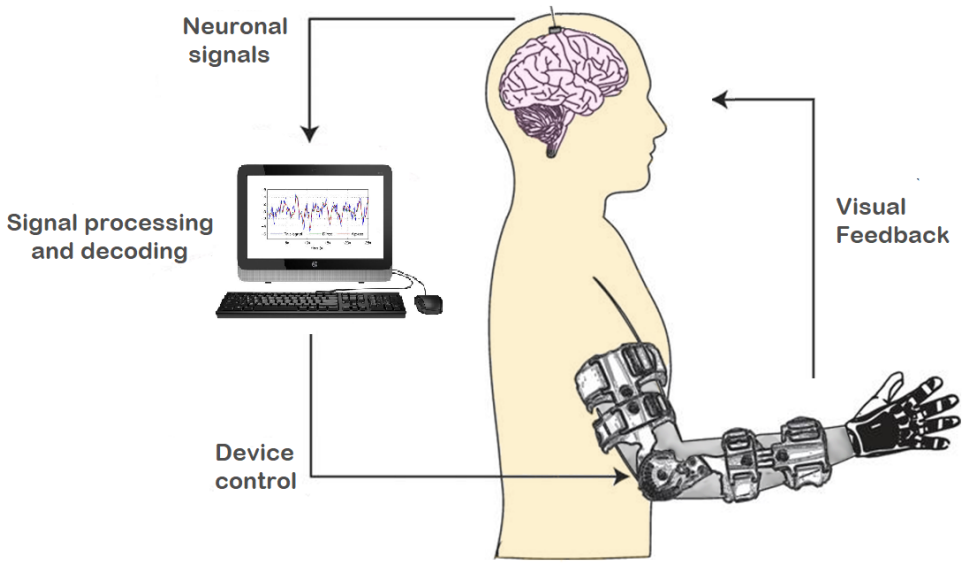

Figure 2.3: Brain-machine interface for motor control. Motor commands are extracted from neuronal activity and sent to external devices to allow real-time neural control. The BMI decoder applies algorithms to the neural input to calculate output variables. The visual feedback loop allows users to revise the performed action. (Adapted from Lopes et al. (2006) (28))

1998 the first implantation of a multi-channel array in primates was successfully accomplished, allowing the simultaneous recording of up to 48 single-neurons across multiple cortical and subcortical regions for several weeks. (1)

Multiple studies were carried out in rodents and monkeys with the intent to understand the physiological properties of the brain, namely plastic adaptation and the learning process that leads to the self-regulation of their activity in response to rewards. As a result, in the early 2000s, it became evident that BMIs were a breakthrough in the field of neuroprosthetic devices, as they have potential to restore mobility to patients with severe paralysis. (1)

Nicolelis et al. (2004) reported in 2004 the extraction of hand movement control from ensembles of subcortical neurons in awake and conscious humans. (31) The same laboratory created in 2009 the first BMI able to decode the kinematics of bipedal walking in rhesus monkeys, and in 2011 the first tool to deliver direct tactile feedback to the subject's somatosensory cortex. (1,32,33)

BMI research yielded several medical applications, from computer-assisted spellers for patients with locked-in syndrome, to exoskeletons that restore bipedal walking, robotic limbs, wheelchairs or avatar bodies. Nonetheless, outside of the clinical context, BMIs applications have emerged in computer gaming, education, and the automobile industry. (1)

2.2.2 Electrophysiological recordings

When designing a BMI, one should determine the most appropriate recording setup. A viable recording device should consider signal longevity, signal stability, digital sampling frequency, the information contained in one single channel, and the invasiveness and risks associated with the setup. High longevity prevents frequent replacement of the device, decreasing the risks associates with surgery complications. Signal stability is important to maximise accuracy. Moreover, the sampling frequency should be minimised to diminish tissue heating and power consumption, while still capturing enough information to control the desired output. (34)

In the history of BMIs, various signal sources have been used to control BMIs, such as near-infrarred spectroscopy of the cerebral blood flow, calcium imaging or magnetoencephalography. (35–37) However, the best compromise between costs, size, and resolution is achieved with electrical signals, making it the most common input sources for BMI motor control. (34)

Electrical signals can be, in descending order of their invasiveness, action potentials (spikes) and local field potentials (LFPs), subdural signals (measured through electrocorticography, ECoG), epidural field potentials (EFPs), and electroencephalographic signals (acquired with electroencephalography, EEG). (34) The less invasive a signal is, the less quality it has, due to the attenuation caused by the skull, scalp and cerebrospinal fluid. (38) It, thus, comes as no surprise that intracortical signals have been proven to contain the most movement-related information among all signals in the context of BMI performance during reaching and grasping. (34)

Intracortical electrical signals can be bandpass-filtered at 0.5 – 300 Hz or 300 – 12.000 Hz to extract, respectively, LFPs and action potentials (spikes). Despite recent studies showing the great deal of information that can be retrieved from LFP, they are yet to outperform action potentials when it comes to decoding kinematics and muscle activity. (34) The action potentials detected by a single electrode can represent the summed neuronal activity around said electrode (MUA) or can be sorted to represent the activity of a single neuron (single-unit activity (SUA)). (34) For a BMI application, Carmena et al (2003) demonstrated that prediction based on the activity of neuronal ensembles (MUA) were considerably more accurate than those based on the activity of individual neurons (SUA). (39)

Long-term multichannel recording of MUA required the development of new designs of implantable arrays. Over the years several options have been created, including microwire recording cubes, Utah array, neurotrophic electrodes, neural dust and endovascular electrodes. However, currently, the Utah array is the only microelectrode implant that is approved by the United States Food and Drug Administration (FDA) for human use, which makes it the most suitable for studies that aim to be reproduced in humans. (1)

Utah arrays have 100 silicon-based rigid microelectrodes, each one ~1.5 mm long, with a

0.4-mm spacing between neighbouring needles, as described in fig. 2.4. The electrode's shafts are coated with polyimide, and their sharpened tips with platinum. The main advantage of this array is that it needs not be attached to the skull. (1) Instead, it may float on the cortex, decreasing the interference that small movements of the brain with respect to the skull cause to recordings over time. (40) On the flip side, the insertion method through a pneumatic gun causes the tissue around the electrodes to scar, precluding these needles from recording neuronal activity endlessly due to fibrous encapsulation. (1)

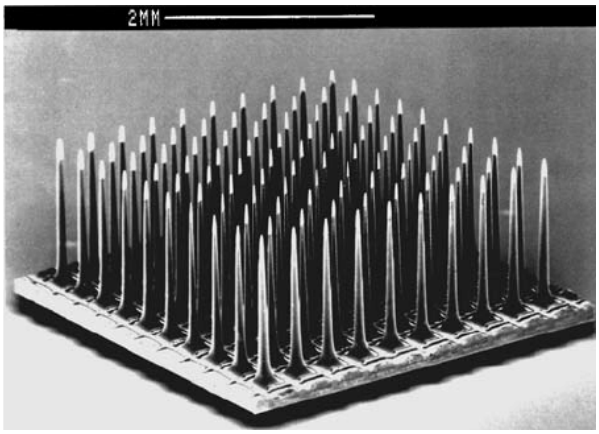

Figure 2.4: The Utah array. A: The 4x4 mm silicon substrate contains 100 penetrating microneedles, each 1.5 mm in length, with an electrode spacing of 0.4 mm. (Reproduced from Kim et al. (2006) (41))

In the context of neuronal recordings, a spike is defined as electrical signals that crosses a predefined voltage threshold. Although this threshold-crossing is prone to confusing artefacts and noise with real neuronal activity, some studies argue that this method has the potential to yield satisfactory results in BMI control. (1, 42)

2.2.3 Information encoding by neurons and neuronal ensembles

Despite almost 60 years of accumulated literature on how neurons encode ethologically-significant information, there is still a critical knowledge gap in how physiological mechanisms operate neuronal circuits. Nonetheless, BMI researchers have sought to extract parameters of interest from neuronal signals, based on empirical evidence showing some degree of correlation between those behavioural variables and neural activity – referred to as neuronal tuning. (1) An example of such was found by Edward Evarts (1966), who described an M1 neuron whose firing increased or decreased when a monkey pulled or pushed a lever. (43)

The consistency of the motor outputs generated by functional BMIs is tightly linked with two basic physiological properties. First, the simultaneous recording from a neuronal ensemble

can compensate for the trial-to-trial variability in the firing rates of individual neurons (neuronal noise). Secondly, neurons can continuously adapt their tuning to new tasks and BMI statistics. (1) Additionally, several principles have been proposed to describe the physiology of neural ensembles, which have highly benefited the development of new projects with BMIs. (44)

The single-neuron insufficiency principle states that single neurons carry very limited information. Adding consecutively more neuronal elements can substantially improve decoding performance. This effect, however, stabilises when the neuronal ensemble reaches a certain size, as describes by the neuronal mass principle. (1,44)

As postulated by the distributed-coding principle, neurons located in multiple cortical areas contribute to the representation of a given behavioural parameter, supporting the belief that information is distributed within the cortex. (1,44) A very pertinent example was described by Matthew Best (2016), showing that both PMd and PMv represent the kinematics of arm reaching and hand grasping. (45)

The neuronal multitasking principle proposes that neurons are not highly-specialised in a task, but rather represent combinations of behavioural parameters. Supporting this idea, the neural degeneracy principle argues that, not only can neuronal ensembles command multiple behavioural outputs, but also different neuronal ensembles can encode the same behavioural output. (1,44) Together, these two principles assure the robustness and flexibility of the neuronal code. (46)

In order to maintain the energy consumption fixed, the conservation of firing principle explains that, if some neurons increase their activity to encode a behavioural parameter, others need to reduce their firing rate proportionally. In fact, the context principle argues that neuronal firing patterns produced in response to an event differ with the circumstances surrounding said event, which shows that the brain is capable of contextualising information. (47)

Neuronal populations alter their properties in response to novel conditions or the learning of new behavioural tasks, as stated by the plasticity principle. Indeed, subjects can only learn to control a BMI and improve their motor performance through cortical plastic adaptations. (1,44) This adaptation translates into an assimilation of the external actuators (e.g. a robotic arm) as if they were an integral part of the subject's body representation, known to exist in our brains. BMI-related cortical plasticity causes changes in both directional tuning and temporal patterns of single neurons (39)

2.2.4 Modelling neuronal activity

Throughout the history of neuroscience research, there has been an undeniable interest in understanding how movement is generated in our brains, which made the motor cortex one

of the most extensively studied cortical areas. Yet, the idiosyncrasies of its neuronal response still lack consistent understanding. (8, 48) Two conflicting models have, since then, proposed possible interpretations of how neuronal responses instruct motion: the representational and the dynamical views.

The representational view explains the activity of single neurons as a function of movement parameters, supporting the premise that the individual motor cortex neurons encode many high-level or abstract movement features according to eq. (2.1). (48)

$$r_n(t) = f_n(param_1(t), param_2(t), param_3(t)...) \tag{2.1}$$

In this equation, $r_n(t)$ is the firing rate of neuron n at time t, and f_n is a function tuning for parameters $param_1(t), param_2(t), param_3(t)...$, such as velocity or target position. (8)

This model was rooted in comprehensive descriptions of single-neuron tuning and was fundamental to take the first steps in the field of brain computation. However, improvements in electrophysiological recording technology have enabled the simultaneous recording of various neurons, opening the doors to develop other models with enough explanatory power to describe more sophisticated aspects of single-neuron and population-level activity. (48)

The dynamical-systems view states that neuronal activity can be predicted from the population activity using a dynamical system, which generates and controls muscle activity directly, as described in eq. (2.2). (48)

$$\dot{\mathbf{r}}(t) = f(\mathbf{r}(t)) + \mathbf{u}(t) \tag{2.2}$$

In this equation, $\mathbf{r}$ describes the population response, or neuronal state, $\dot{\mathbf{r}}$ is its temporal derivative, f is an unknown function and $\dot{\mathbf{u}}$ is some external input. (8)

Inspired by their observations that rhythmic muscle contraction in both a swimming leech and a walking monkey matched the oscillation frequency of cortical responses (figs. 2.5(a) and 2.5(b)) Churchland et al. (2012) proposed treating neuronal responses as a dynamical system itself. They concluded that, if individual neurons have oscillatory activity, the overall population response should rotate with time. (8) In order to test this hypothesis, this team developed jPCA - *see* chapter 3 for a detailed description -, a method that seeks linear combinations of principal components (obtained via PCA) that capture the rotational structure in a population of neurons. (8) Indeed, quasi-oscillatory neuronal responses not only were found in a motor neuronal population during a reach-to-grasp task (fig. 2.5(c)), but they capture a large portion of the variance observed. (8, 48) This may come as a surprise, as the task itself is not rhythmic, but the rotational patterns were still present. (8)

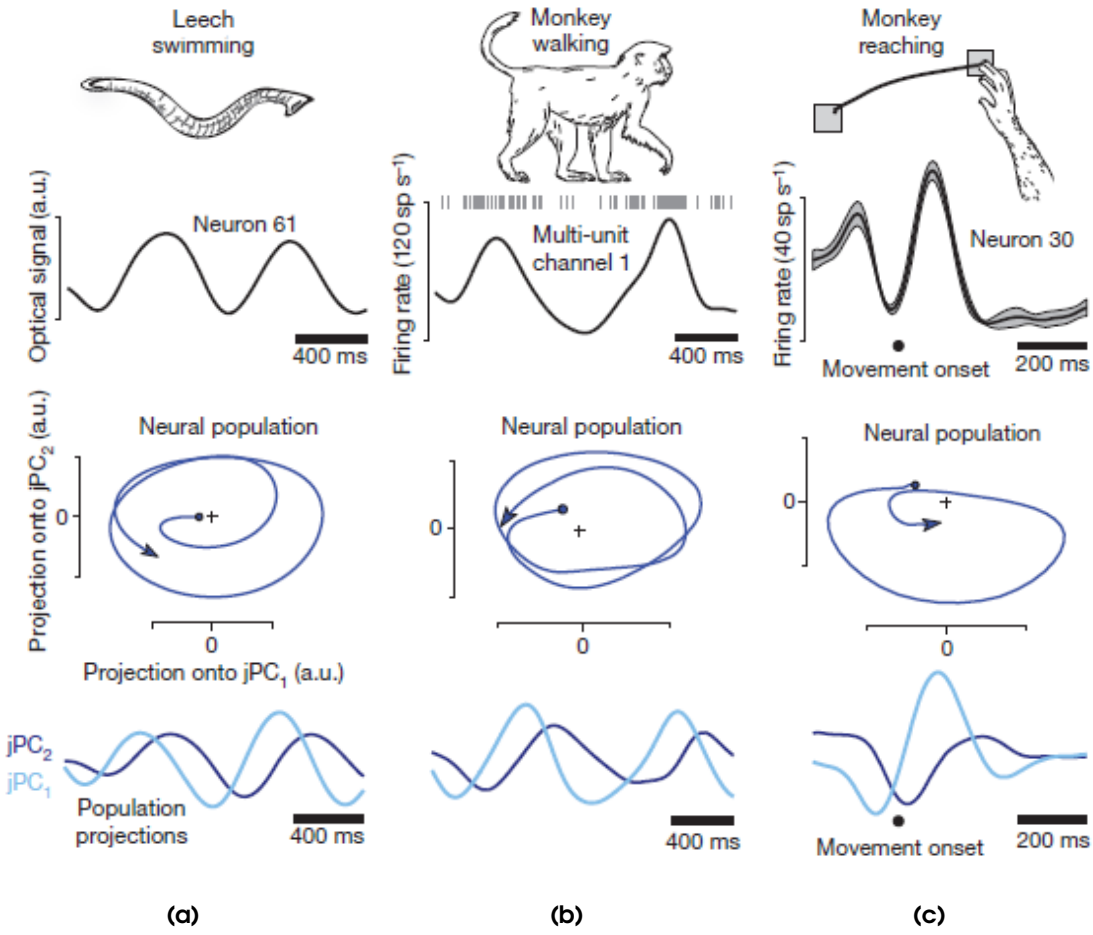

Figure 2.5: Oscillation of neural firing rates during three movement types. **(a)** Response of 1 of 164 neurons in the isolated leech central nervous system during a swimming motor pattern (top), and projection of the leech population response on the first jPCA plane. **(b)** Same as in **(a)** for the MUA of 1 out of 96 electrodes implanted in the arm representation of the caudal premotor cortex of a monkey. **(c)** Same as in **(a)** for the response of 1 out of 118 electrodes recorded from the monkey's motor cortex. (Reproduced from Churchland et al. (2012) (8))

Since then, many studies report that the rotational patterns underlying motor cortex trajectories during reach-to grasp tasks can describe the patterns of motor activity. (2,8,49) In particular, muscle activity should be a direct combination of neuronal population rotations. (48)

2.2.5 Decoding of neuronal signals

There have been developed numerous statistical and machine-learning methods for transforming neuronal activity into interpretable signals for the artificial actuators of BMIs. Be it for online (also called real-time) or offline decoding, these BMI decoding algorithms have to be prepared to receive multiple inputs from the neuronal recording channels and output the corresponding

behavioural variables. (1)

Decoding algorithms consist of a set of independent parameters that map to a small set of output variables. Setting these parameters involves what is called decoder training. (1) In neuroscience experiments, it is common practise to perform a ten-fold cross validation to evaluate decoding performance: 90% of the trials are used to train the decoder, constituting the training set, and another 10% of the trials are evaluated against their real labels, forming the test set. (50) When the decoder reaches high performance, the BMI mode of operation can begin, which in practise means that the model can perform online decoding. (1)

In real-life applications, when decoding of fine movements and continuous space, there are some decoders that are more adequate than others. Namely, the Wiener filter produces an estimate of the target neuronal population activity, that is used to extract limb kinematics and other parameters. (1) Indeed, Carmen et al. (2003) used a Wiener filter to generate the x and y components of movement. (39) The Kalman filter is another popular algorithm that issues motor commands according to the current state of the motor system and neuronal activity. (51) The point-process model is an analogue of the Kalman filter, that estimates the probability of a neuron to produce a spike based on its spiking history and population activity. Another increasingly popular option is to use recurrent neural networks, a dynamical artificial neural network that can interpret temporal dynamic behaviour. (1)

However, for simpler problems with discrete choices, discrete algorithms are equally popular. They yield equally good outcomes, without the need of extensive datasets for training. In this class, we can include linear discriminant analysis (LDA), SVM, multilayer perceptron, artificial neural networks, hidden Markov models, k-nearest neighbours, and nonlinear Bayesian classifiers. (1)

In the context of BMI applications, good decoding performance has been reported for SVMs. (52–55) Furthermore, Wang et al. (2012) argue that SVMs, as a wide-margin classifier, are suited to non-Gaussian data, and that its performance is qualitatively and quantitatively better in cross-validation than other decoders, making this method appropriate for decoding movement and timing of movements. (55)

Chapter 3

Methods

This section includes the description of the protocols, materials and methodologies needed to reproduce the experiments and analyses in this thesis.

3.1 Neuronal data acquisition

3.1.1 Electrophysiological recordings

The neuronal data used in the present work was recorded in two distinct periods from one male rhesus monkey (*Macaca mulatta*, 8 kg). The data acquired in the first and second experiments are analysed in parts 2 and 1 of the results, respectively.

In both instances, a titanium head post was fixed to the subject's skull with dental acrylic and ceramic screws, and 96-channel micro-electrode Utah arrays with 1.5-mm-long electrodes and an electrode spacing of 0.4 mm (Blackrock Microsystems, UT, USA) were implanted on the premotor cortex.

In the first experiment, a single Utah array was positioned in the dorsal premotor area F2, PMd. In the second experiment, two Utah arrays were inserted in the dorsal premotor area F2, PMd, and the ventral premotor area F5c, PMv. In both cases, the electrodes were implanted in contralateral side to the subject's working hand – the left hand – with the guidance of stereotactic coordinates and anatomical landmarks (the principal and arcuate sulci, in fig. 3.1(b)). The positioning of the arrays was verified with anatomical fMRI scans (fig. 3.1(a)).

The monkey was kept under general anaesthesia with propofol (10 mg/kg/h) and strict aseptic conditions, during all surgical procedures.

In both experiments, all channels were connected to digital headstages (Cereplex M, Blackrock Microsystems, UT, USA). In the first experiment, all 96 channels were connected to a 96-

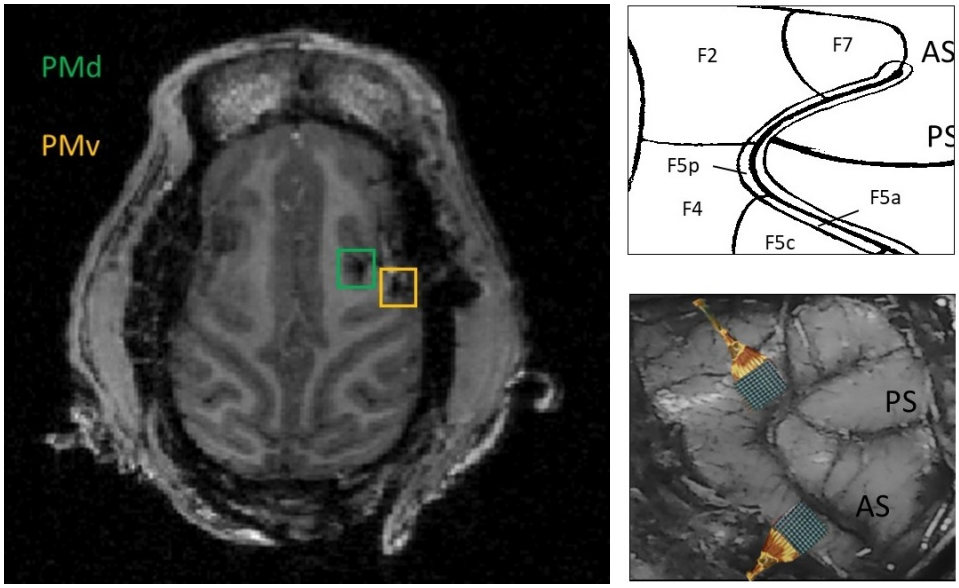

Figure 3.1: Location of the chronically implanted Utah arrays in the second experiment. **(a)** Transversal structural magnetic resonance scan of the subject's brain showing the location of the dorsal premotor (PMd) area F2 (green square) and ventral premotor (PMv) area F5c (yellow square). **(b)** Schema of the premotor areas and picture of the exact locations of implantation. PS: Principal sulcus; AS: Arcuate sulcus.

channel digital neural processor (Blackrock Microsystems, UT, USA); in the second experiment, only two banks of 32 channels from each headstage were connected to a 128-channel digital processor (Blackrock Microsystems, UT, USA), according to table 3.1. The recordings are then sent to a Cerebus data acquisition system (Blackrock Microsystems, UT, USA) at a 30 kHz sampling frequency.

Table 3.1: Channels of the Utah arrays recorded in PMd and PMv of monkey J in the second experiment.

Dataset	PMd	PMv
monkey J1	1-64	1-64
monkey J2	33-96	33-96
monkey J3	1-32, 65-96	1-32, 65-96
monkey J4	1-64	1-64
monkey J5	33-96	33-96
monkey J6	1-32, 65-96	1-32, 65-96

A multiunit detection trigger was set to detect spikes in the neural activity. The threshold was set at 95% of the maximum noise recorded for each individual channel, and every time this

value was exceeded the time instant was saved.

All medical and experimental protocols were approved by the ethical committee on animal experiments of KU Leuven and performed according to the *National Institute of Health's Guide for the Care and Use of Laboratory Animals* and the EU Directive 2010/63/EU.

3.1.2 Experimental paradigm

During experimental procedures, neuronal activity was recorded while the monkey sat upright in a primate chair with its head fixed. A custom-built object (15 cm diameter) containing three identical spheres (2.5 cm of diameter) was placed 34 cm in front of the monkey. The small spheres were attached to the large object with springs, allowing to be pulled. Each sphere contained a blue LED, used to signal the object to be grasped, positioned in an angle of 120º relative to the other two spheres. The larger object incorporated a green LED in the centre to indicate the go cue.

In each task repetition (or trial), the monkey was instructed to reach for and grasp one of the spheres in a pseudo-random order, while neuronal activity was recorded. Each sphere corresponds to a different experimental condition, as all spheres require a different reaching direction.

To start a trial, the monkey positions its left hand (contralateral to the implanted motor areas) on a hand-rest in complete darkness. The movement of the hand is monitored with an IR laser beam, which is interrupted when the hand is in the resting position. Shortly after the start, the green LED goes on and, after at least 500 ms of fixation on the green light, one blue LED goes on (object onset), indicating which of the sphere the monkey has to pull. Simultaneously to the blue LED, an external light source illuminates the object from above. After a variable time, the green LED is turned off (go cue), instructing the monkey to reach for the cued object. The time interval between the instant when the monkey effectively releases the hand-rest (lift hand) and pulls the sphere (pull object) cannot exceed 1000 ms, otherwise the trial is considered incorrect. A correctly-executed trial is rewarded with juice. The inter-trial interval is 1000 ms. Figure 3.2 illustrates four distinct moments in a trial.

The experimental setup described was the same for all datasets utilised in this work.

3.1.3 Data Processing

The recorded data correspond to spike times of the MUA detected by each electrode. To obtain the spike rates at a time point, the number of recorded spikes in a 20-ms bin centred in that instant is counted and divided by the duration of the interval - hence the unit of spike/s. The

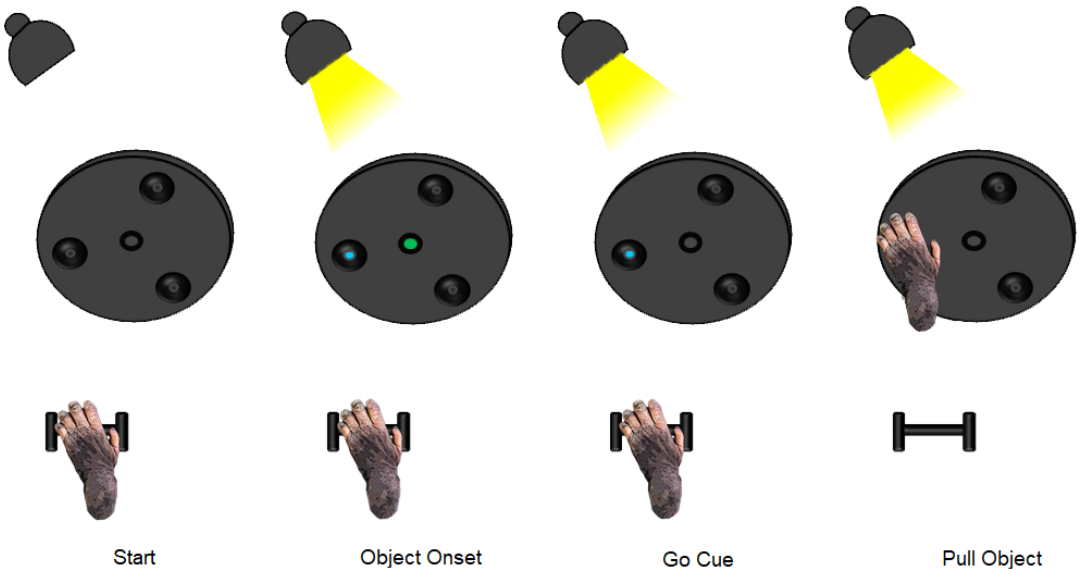

Figure 3.2: Schema of the delayed reach-to-grasp task. The monkey was trained to start a trial from a resting position, and grasp the sphere with the blue LED after the go cue (dimming of the green LED). The object was illuminated during the duration of the trial.

rate is calculated every 10 ms, and the results are plotted in Peristimulus Time Histograms (PSTHs).

The signal in each channel is normalised by subtraction of the average spike rate in a 300-ms interval recorded immediately before object onset - baseline activity. Additionally, since it is mathematically not possible to calculate the spike rates with an infinite precision, the PSTHs are smoothed with a 24-ms Gaussian filter, in order to simulate temporal continuity.

3.2 jPCA

A major concern in neuroscientific experiments is organisation and visualisation of data. Each experiment may record from many channels, and the signal recorded in a channel varies with time and experimental condition. As a result, this will create a high-dimensional state-space, which is hard to work with.

With the goal of facilitating structural exploration of neuronal data, in 2010, the Churchland laboratory developed jPCA, a dynamical variant of Principal Componenent Analysis (PCA). The code for jPCA is available at `https://churchland.zuckermaninstitute.columbia.edu/`. While PCA finds direction of maximal variance, jPCA finds planes of significant rotational structure. (8) Intuitively, one may argue that the plane with the strongest rotations could, in practice, capture very little variance. Thereby, PCA was always applied before jPCA, ensuring that the rotational structure found could explain a considerable amount of data variance.

Data dimensionality was consistently reduced to 6 principal components for both experi-

ments using PCA. This choice is based on previous papers by Churchland *et al.* (2012), where the group mentions that, for a 96-dimensional state-space, 6 principal components account for $50 - 70\%$ of data variance. (8)

Although completely self-reliant and independent from one another, both methods play important and complementary roles in this project. Appendix A includes a brief example on how PCA and jPCA ares used, along with a geometrical interpretation of jPCA.

In between processing stages, data is normalised. For each of the 6 projections, the firing rate range across all conditions and times is calculated and each entry in that projection is divided by that value plus five, e.g. a neuron with FR range of 5 gets mapped to a range of 0.5. Afterwards, the across-condition mean from each projection is subtracted to the activity in each condition. This last step assures that we are modelling only the differences between conditions.

3.2.1 Mathematical description

Considering a matrix $X = [X_1\ X_2\ ...\ X_c]^T \in \mathbb{R}^{ct \times n}$ that represents the whole set of time series neuronal data, where $X_j \in \mathbb{R}^{t \times n}$ is the time series neuronal data for condition $j = 1, ..c$, t is the number of time points of neuronal data analysed for a single condition, c is the number of conditions, and n is the dimension of the neuronal states (6, after PCA).

Since the goal is to describe the temporal dependency of states, we are working in the field of dynamics. The simplest dynamical system that can fit to data is a time-invariant linear dynamical system in the form of $\dot{X} = XM$ for $M \in \mathbb{R}^{n \times n}$. The notation adopted in this project establishes $\dot{X} = [\dot{X}_1\ \dot{X}_2\ ...\ \dot{X}_c]^T$ with $\dot{X}_j = X_j(2:t,:) - X_j(1:t-1,:)$ and X is, in reality, $X(1:t-1,:)$, where the index $(a:b,:)$ indicates all columns from rows a to b. This way, all dimensions are coherent.

Explain

3.2.1.A Solving for the generic linear dynamical system

Linear dynamical systems with unconstrained summary matrices, like M, describe the whole structure of data - both expansions/contractions and rotations - indiscriminately.

Solving $\dot{X} = XM$ for M can be reduced to the simple least squares problem in eq. (3.1), where we want to find the system M that minimises the error between $\dot{X}$ and XM. In this expression, $\mathrm{tr}(.)$ denotes the trace of a matrix, and $\|.\|_F$ is the Frobenius norm.

I'm happy to see you have used part of the mathematical formulation we addressed in SSB

$$\begin{aligned} M^* &= \underset{M\in\mathbb{R}^{n\times n}}{\operatorname{argmin}} \sum_{r=1}^{c(t-1)} \sum_{p=1}^{n} (\dot{X}(r,p) - X(r,:)M(:,p))^2 \\ &= \underset{M\in\mathbb{R}^{n\times n}}{\operatorname{argmin}} \left\| \dot{X} - XM \right\|_F \\ &= \underset{M\in\mathbb{R}^{n\times n}}{\operatorname{argmin}} \operatorname{tr}((\dot{X} - XM)^T(\dot{X} - XM)) \end{aligned} \tag{3.1}$$

The closed-form solution to this minimisation problem, in eq. (3.2), can be derived applying the properties of matrix-by-matrix differentiation to $\nabla_M \operatorname{tr}((\dot{X} - XM)^T(\dot{X} - XM)) = 0$.

$$M^* = (X^T X)^{-1} X^T \dot{X} \tag{3.2}$$

Although this is not exactly the goal of jPCA, this system will play an important role in the model's quality assessment.

3.2.1.B Solving for the constrained linear dynamical system

jPCA is particularly interested in finding planes of significant rotational dynamics. It is, therefore, necessary to find a system that describes these rotations and disregards expansions/contractions in state-space.

Every linear transformation, like M, can be decomposed into a symmetric and a skew-symmetric part, satisfying $M = M_{sym} + M_{skew}$. The symmetric part of M, defined as $M_{sym} = M_{sym}^T = (M + M^T)/2$, has real eigenvalues and eigenvectors and, therefore, describes expansions and contractions of data. The skew-symmetric part of M, which satisfies $M_{skew} = -M_{skew}^T = (M - M^T)/2$, has complex eigenvectors and zero-valued or purely complex eigenvalues (in conjugate pairs), thus accounts for rotations in the data. A geometrical interpretation for both dynamical systems is included in chapter 2.

The system $\dot{X} = XM$ must now be solved for $M \in \mathbb{S}^{n\times n}$, where $\mathbb{S}^{n\times n}$ is the set of $n \times n$ skew-symmetric matrices.

Similarly to the unconstrained case, M_{skew} can be determined using a least squares approach with a slightly heftier notation by rewriting the original problem as a vector problem. Introducing the vector $\mathbf{m} = M(:) \in \mathbb{R}^{n^2}$, corresponding to the unrolled matrix $M \in \mathbb{R}^{n\times n}$, the vector $\dot{\mathbf{x}} = \dot{X}(:) \in \mathbb{R}^{c(t-1)n}$, and the matrix $\tilde{X} \in \mathbb{R}^{c(t-1)n\times n^2}$, a block diagonal matrix with X repeated n times.

The key step to restrict the set of possible solutions is to represent the skew-symmetric matrix M_{skew} as a vector of $n(n-1)/2$ free parameters, rather than n^2. Thus, we can specify $\mathbf{k} \in \mathbb{R}^{n(n-1)/2}$, and the necessary linear transformation, $H \in \mathbb{R}^{n^2\times n(n-1)/2}$ to map from this vector

onto the space of $\mathbf{m} \in \mathbb{R}^{n^2}$. The role of H is to take an element of $\mathbf{k}$, kr, place that value and its symmetric in the element of $\mathbf{m}$ corresponding to the entries (i, j) and (j, i) of M_{skew}.

The eq. (3.3) displays the 3 equivalent least squares problems described.

$$\begin{aligned} M^* &= \underset{M \in \mathbb{S}^{n \times n}}{\operatorname{argmin}} \left\| \dot{X} - XM \right\|_F \\ m^* &= \underset{m \in \mathbb{R}^{n^2}}{\operatorname{argmin}} \left\| \dot{\mathbf{x}} - \tilde{X}\mathbf{m} \right\|_2 \\ k^* &= \underset{k \in \mathbb{R}^{n(n-1)/2}}{\operatorname{argmin}} \left\| \dot{\mathbf{x}} - \tilde{X}H\mathbf{k} \right\|_2 \end{aligned} \tag{3.3}$$

Following on from the unconstrained problem, the closed-form expression of the solution to the constrained problem can be constructed by grouping H with $\tilde{X}$, as in eq. (3.4).

$$k^* = (H^T X^T X H)^{-1} H^T X^T \dot{\mathbf{x}} \tag{3.4}$$

The summary matrix, M_{skew}, that describes rotational dynamics in data can now be constructed. As previously mentioned, the eigendecomposition of this matrix creates a ranked set of orthonormal vectors that come in complex conjugate pairs. To visualise the data structure on a plane defined by these conjugate pairs, an equivalent pair of real-valued vectors, $\{\mathbf{u}_1, \mathbf{u}_2\}$, has to be found. For that purpose, each complex conjugate pair of vectors $\{\mathbf{v}_1, \mathbf{v}_2\}$ is combined as: $\mathbf{u}_1 = (\mathbf{v}_1 + \mathbf{v}_2)/\sqrt{(2)}$, $\mathbf{u}_2 = j(\mathbf{v}_1 - \mathbf{v}_2)/\sqrt{(2)}$, where the eigenvalue of $\mathbf{v}_1$ has a positive imaginary part. This new basis is then rotated so that the data is mainly spread along the horizontal axis and most of the trajectory displayed rotates anticlockwise. The projection of the data onto this plane corresponds to the real and imaginary parts of the projection onto the eigenvector with the positive eigenvalue. Discarding the imaginary component of the other eigenvector does not lead to any loss of information, because the projections of the data onto each eigenvector of a conjugate pair are exactly symmetric with respect to the real axis. For a more geometrical explanation of the method, the reader can resort to appendix A.

Since we are working in a 6-dimensional space, the eigendecomposition of M_{skew} will find three pairs of complex conjugate eigenvector, and, consequently, three orthogonal planes. These planes are numbered in decreasing order of the magnitude of the (purely imaginary) eigenvalues corresponding to the eigenvectors that define them, meaning that the first plane captures the most rotational structure.

In this work, all three projection onto the jPCA planes are presented.

3.2.2 Variance captured by a jPCA plane

It is only natural that a 2-dimensional surface cannot represent the entirety of the high-dimensionality data variance. Unless data only varies along 2 dimensions, there must always be some loss of information inherent to the data projection.

As the two vectors that define a plane are orthogonal, $\mathbf{u}_1 \perp \mathbf{u}_2$, the data variance explained by the plane they form is the sum of the individual variances captured by each projection, determined as in eq. (3.5).

$$\begin{aligned} VarCapjPCplane(p) =& VarCapjPC(2p-1) + VarCapjPC(2p) \\ VarCapjPC(i) =& \frac{\sum_{r=1}^{ct} Xproj(r,i)^2}{\sum_{r=1}^{ct}\sum_{b=1}^{L}(R(r,b)-\overline{R})^2} \end{aligned} \tag{3.5}$$

In this expression, $Xproj = X\ [\mathbf{u}_1\mathbf{u}_2] \in \mathbb{R}^{ct\times 2}$ is the projection of X (after PCA) onto a jPCA plane $[\mathbf{u}_1\mathbf{u}_2]$, $R \in \mathbb{R}^{ct\times E}$ is the processed spike rates (before PCA) for each of the recorded channels (64 or 96 depending on the the experiment), and $\overline{R} = \frac{1}{ct}\sum_{r=1}^{ct}\sum_{b=1}^{L} R(r,b)$ is the average spike rate after processing.

3.2.3 Goodness of fit

The quality of the fit of the constrained model, M_{skew}, can be assessed using the coefficient of determination (CD), R^2, calculated as in eq. (3.6) (note that $\dot{\mathbf{x}} = \dot{X}(:)$ is mean-centered after PCA). This measure evaluates the model's ability to predict $\dot{X}$ from X.

$$R^2_{M_{skew}} = \frac{\sum_{b=1}^{c(t-1)n} \dot{\mathbf{x}}(b)^2 - \boldsymbol{\epsilon}_{M_{skew}}(b)^2}{\sum_{b=1}^{c(t-1)n} \dot{\mathbf{x}}(b)^2}, \qquad \boldsymbol{\epsilon}_{M_{skew}} = \dot{\mathbf{x}} - \tilde{X}H\mathbf{k} \tag{3.6}$$

Following the work of Jonathan A. Michaels *et al.* (2016) and Michaels et al. (2016), the ratio between the coefficient of determination of the constrained system, M_{skew}, and the unconstrained system, M, will also be used as a statistical indicator of the goodness of fit. (48) The rotational goodness-of-fit ratio (RGR), as calculated in eq. (3.7), indicates how much of the data variance captured by general linear dynamics are explained by purely rotational dynamics or, in other words, how circular the trajectory is. (8) If linear dynamics are dominated by rotations, M will naturally be very similar to M_{skew}, and M_{skew} will fit data nearly as well as

M. (56)

$$RGR = \frac{\sum_{b=1}^{c(t-1)n} \dot{\mathbf{x}}(b)^2 - \boldsymbol{\epsilon}_{M_{skew}}(b)}{\sum_{b=1}^{c(t-1)n} \dot{\mathbf{x}}(b)^2 - \boldsymbol{\epsilon}_{M}(b)} = \frac{R^2_{M_{skew}}}{R^2_{M}}, \qquad \begin{aligned} \boldsymbol{\epsilon}_{M_{skew}} &= \dot{\mathbf{x}} - \tilde{X}H\mathbf{k} \\ \boldsymbol{\epsilon}_{M} &= \dot{\mathbf{x}} - \tilde{X}\mathbf{m} \end{aligned} \tag{3.7}$$

3.3 Support vector machine

The decoding of movement direction (task condition) and movement onset is performed with the MATLAB toolbox for SVM.

As a supervised learning algorithm, SVM training requires a set of labelled observations, each one characterised by various features. The set of observations is divided into two sets: a training set, that contains about 90% of the observations and is used to train the model, and a test set, constituted by the remaining 10%, that evaluate the performance of the new-found model.

Firstly, the feature space is significantly reduced using the MATLAB function for feature selection, *sequentialfs*. This function selects a subset of features that best predict the label of the observations. Starting from an empty subset, only the feature with best results in a 10-fold cross validation procedure is added to that subset. If the new subset does not have a better performance than the previous subset, no new feature is added. (57)

To train the algorithm, data dimensionality is raised using a kernel. This kernel trick improves the resemblance of the input to linearly separable data. The *radial basis function* kernel was empirically found most suitable for the analyses developed.

For each model, Mdl, trained with the selected features, the labels of the test set are predicted and the accuracy is determined as the percentage of correct label predictions.

For statistical significance, each set of decoding conditions (decoding intervals and feature design) is repeated 100 times.

3.3.1 Decoding movement direction

Since SVMs is a binary classifier, and there are three conditions to classify, the problem is approached with an error-correcting output code (ECOC) model with an SVM learner template. (58) This multiclass classifier was implemented using the MATLAB function *fitcecoc* with a *onevsone* coding design. (59)

In order to improve performance, some hyperparameters, such as the kernel scale and the box constraint, are optimised each time the machine is trained. The kernel scale defines the

adequate *radial basis function* width to cluster elements of the same class. The box constraint applies a cost to the misclassification of data, which is higher the larger the constraint. (59)

For each dataset and each interval, each trial was considered an observation and jPCA was performed to find the six jPCA projections that captured the most rotational structure among all trials. For each observation, each instant in a jPCA projection was considered a feature. For instance, if jPCA analysis is performed on a sequence of 10 time points and 6 projections are outputted, each observation will have $10 \times 6 = 60$ features.

Similarly, when decoding with all neurons instead of jPCA projections, each time instant (within the interval of the analysis) in each channel is acknowledged as a feature. Following the example above, for 96 neurons, each observation has $10 \times 96 = 960$ features.

3.3.2 Decoding movement onset

As movement onset is a binary problem - motion is either imminent or not -, this decoding problem was solved by directly fitting an SVM model to the data. In this circumstance, the MATLAB function *fitcsvm* was used. (60)

Similar to the procedure to decoding movement direction, the optimal kernel scale and box constraint are found each time the SVM is trained.

For this application, each trial retrieves two observations: one sampled in the 300 ms immediately before movement onset (labelled 'Onset'), and another one sampled from the interval (50; 350) after object onset (labelled 'Not Onset'). The two groups are joined, forming a set of observations with double the number of trials in the dataset. Next, the group of observations is subject to jPCA analysis.

The feature design, with jPCA or all channels, is similar to the analogous procedure described for decoding movement direction.

3.4 Material Specifications

All data were analyzed using custom written Matlab scripts (MATLAB R2018b, Mathworks, MA, USA). The Matlab's Java Heap Memory preferences were set to 4.014 MB (especially important for the decoding part), and the scripts were run on a Precision 5820 Tower (Dell, TX, USA) with a 12-core 3.6 GHz processor, 16384 MB RAM and SSD.

Chapter 4

Results

This chapter consists of three parts with a complementary purpose. First, the spikes are processed to represent the neuronal activity in the form of PSTHs. Secondly, jPCA is used to find rotational structure in the state-space representation of the neuronal activity in PMd and PMv. Finally, jPCA projections are used to decode movement direction and movement onset.

4.1 Neuronal activity in PMd and PMv

This section illustrates the processing stages of the neuronal activity recorded, and analyses the patterns of activity in PMd and PMv during a reach-to-grasp task. For this application, datasets J1-6 in table 3.1 are used.

4.1.1 From spikes to spike rates

The classical view of neural computation states that information is encoded in the spike rate of neurons. It is also widely believed that the temporal coordination of spikes plays an important role in information encoding. (61) Spike rates are, however, a mathematical construction on spike trains, and are, therefore, not directly observable. The spike times, on the other hand, are recorded by the Utah array.

Figure 4.1 represents an example of a spike train recorded in one channel for several trials. Each channel records MUA, meaning that many neurons are detected in each channel. As the time interval between the events of interest differs trialwise, this plot displays blank regions with the sole purpose of aligning these events and facilitate reading. This figure illustrates a channel that detects a higher temporal density of spikes starting after the monkey sees the object (Object) until pulling it (Grab).

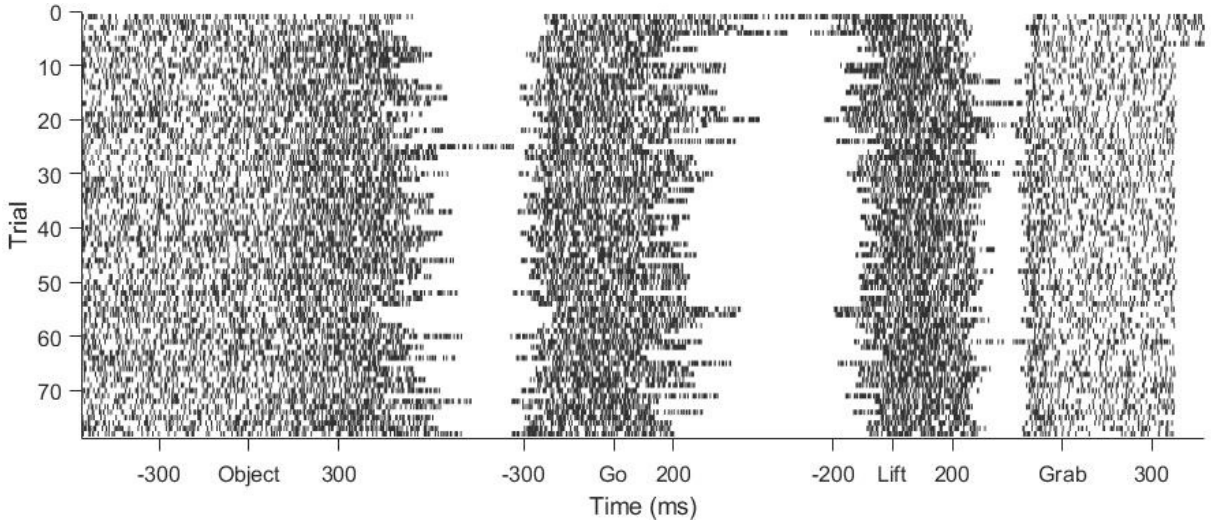

Figure 4.1: Temporal raster plot of the spikes detected in channel 16 of the Utah array implanted in PMd in monkey J3-PMd. Each row represents a trial of a reach-to-grasp task, and each black line symbolises the binary existence of spikes in 1-millisecond bins. The trials are aligned to four events of interest.

From these spike times, it is possible to calculate the rates of activity in the form of histogram of spike counts per interval of time (*see* chapter 3). These so called PSTHs are a more objective way of looking at both the rate and the timing of neuronal spikes, compared to raster plots. To approximate the standard response of a channel (or group of neurons), one can align several trials to an event and compute the average response.

Figure 4.2(a) represents the PSTH of the same channel in fig. 4.1 averaged over trials. This plot shows it in a clearer fashion, that the neuronal activity in this channel is, indeed, considerably more prominent between object onset and grab object. Furthermore, PSTHs facilitate a relative comparison of activity between time bins, which was not possible with a raster plot.

To improve the resemblance to a time continuous signal, the resuting PSTHs are frequently smoothed. In the present work, a 24-ms Gaussian filter was used to obtain the PSTH in fig. 4.2(b).

4.1.2 Common response profiles

The across-channel average PSTH of a channel paints a non-realistic, but nonetheless useful, picture of how the neurons in that area respond to a stimulus on average. It is worth noting that averaging across channel may attenuate some important responses of individual channels, so a more thorough analysis is not dispensable.

A channel, or cluster of neurons in the case of multiunit activity, is stimulus-responsive when its activity relatively to the baseline changes after the subject is exposed to stimuli - object visualisation or go cue, in the context of this work. A channel is direction-selective, if its activity

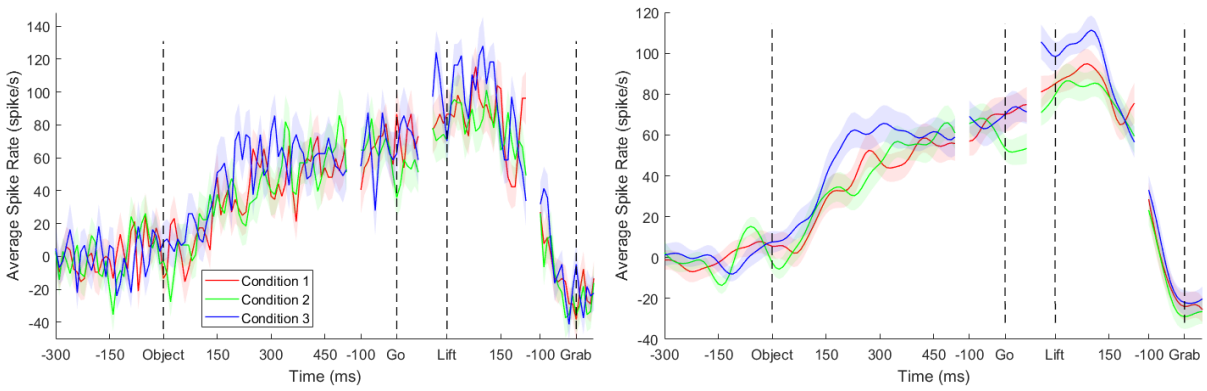

(a) **(b)**

Figure 4.2: **(a)** Unfiltered PSTH of the multi-unit activity detected in channel 16 of the array implanted in PMd of monkey J3 during a reach-to-grasp task. The trials are aligned to four events of interest, and each trace plots the average spike rate across trials for one experimental condition. The standard error is plotted in a lighter shade around the PSTH for each condition. **(b)** The same as (a), filtered with a 24-millisecond Gaussian kernel.

is tuned by the position of the object.

According to fig. 4.3(a), one may expect to observe responsiveness in PMd at almost any time during the reach-to-grasp task; however, direction-selectivity should be more prominent from around 300 ms after the object onset until approximately 150 ms after the monkey lifts the hand. A very different behaviour is predicted for PMv, as the average activity of the recorded channels shows more responsiveness and selectivity after the monkey lifts the hand, as shown in fig. 4.3(b).

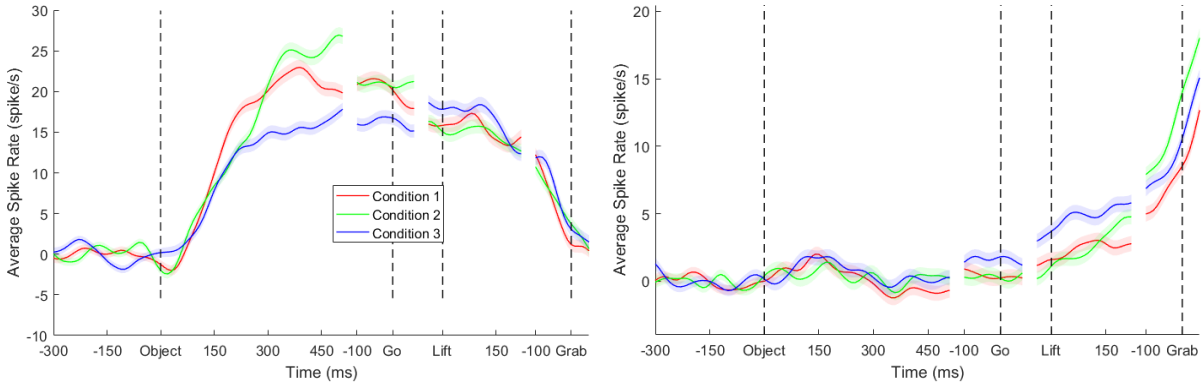

(a) **(b)**

Figure 4.3: Across-channel average PSTH of the multi-unit activity recorded in all 64 channels in **(a)** PMd and **(b)** PMv, for monkey J3. The trials are aligned to four events of interest, and each trace plots the average spike rate across trials for one experimental condition. The standard error is plotted in a lighter shade around the PSTH for each condition.

In both cases, it is not expected to observe responsiveness before object onset, as there is no stimulus. But, even after object onset, the average behaviour is far from describing accurately the nature of single-channel activity. In fact, a variety of response profiles can be observed in both areas during reach-to-grasp tasks.

In PMd, some channels are more responsive during the visual phase, with a steep increase in activity after object onset followed by a return to the baseline when the monkey is instructed to move its hand (fig. 4.4(a)). Other channels are more responsive during the movement phase, and, therefore, have a crescent activity that reaches its peak after the monkey lifts its hand (fig. 4.4(b)). There can also be channels with strong activity during both visual and the movement phases (fig. 4.4(c)). It is also common to record channels with strong multi-unit activity around the moment when the monkey pulls the object (fig. 4.4(d)), which makes these chan-

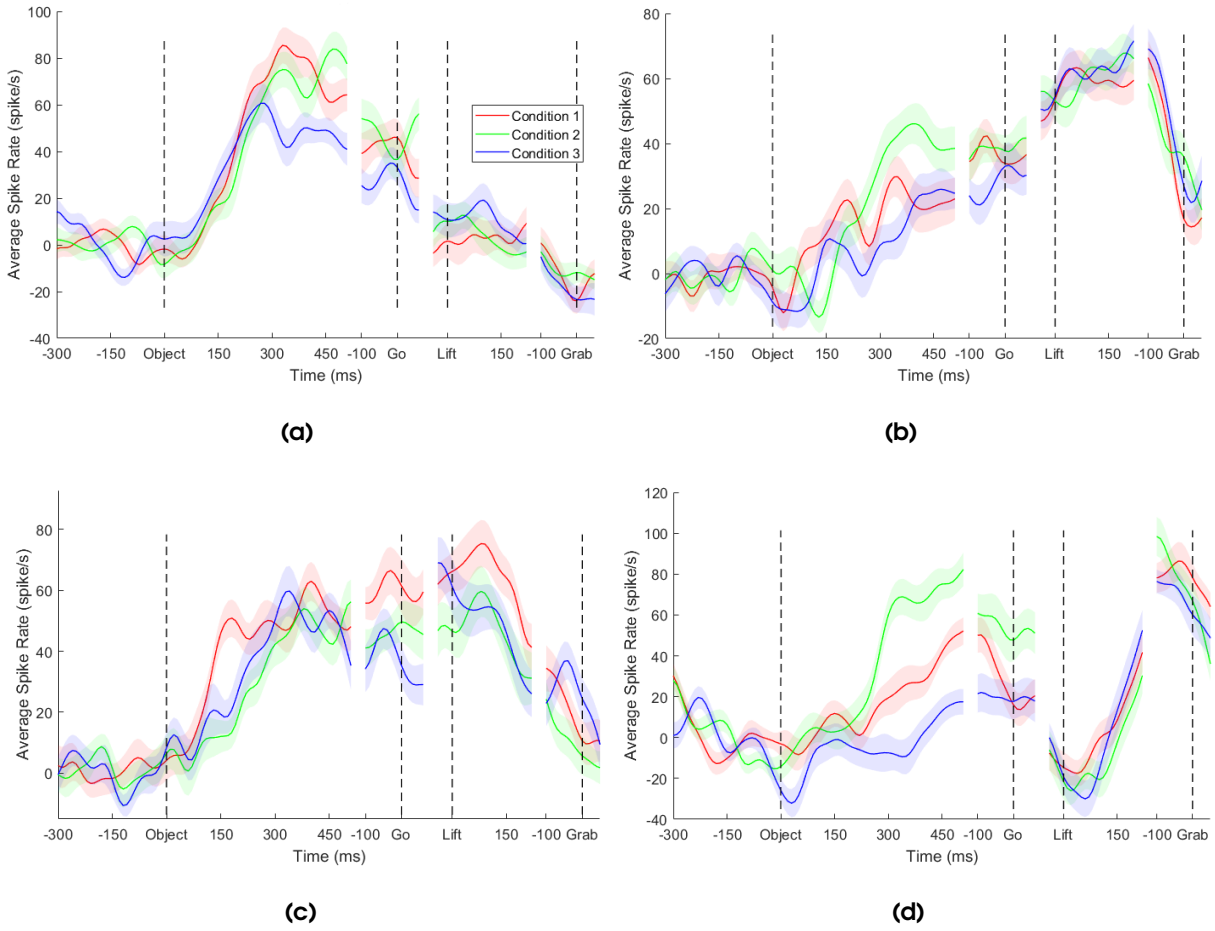

Figure 4.4: Filtered PSTH of the multi-unit activity detected in PMd of monkey J3 during a reach-to-grasp task: **(a)** channel 10, **(b)** channel 23, **(c)** channel 27 and **(d)** channel 15. The trials are aligned to four events of interest, and each trace plots the average spike rate across trials for one experimental condition. The standard error is plotted in a lighter shade around the PSTH for each condition.

nels more responsive to that event.

Most of the across-trial average multi-unit activity recorded in each channel is not equally responsive to all conditions, which is clear in all PSTHs in fig. 4.4. This high selectivity is especially marked during the visual phase. It is also worth mentioning that there are also several channels that are non-responsive to these stimuli, which translates into their spike rate fluctuating around the baseline throughout the whole trial.

In PMv, with some rare exceptions of channels with a strong visual response, most multi-unit activity was recorded during the movement phase, after the monkey lifted the hand. However, contrarily to what one may perceive from fig. 4.3(b), some channels are selective to direction right after object onset (fig. 4.5(a)), and others show more selectivity after the movement starts (fig. 4.5(b)). Similar to PMd, there are also some channels that do not respond (at least significantly) to the stimuli in these tasks.

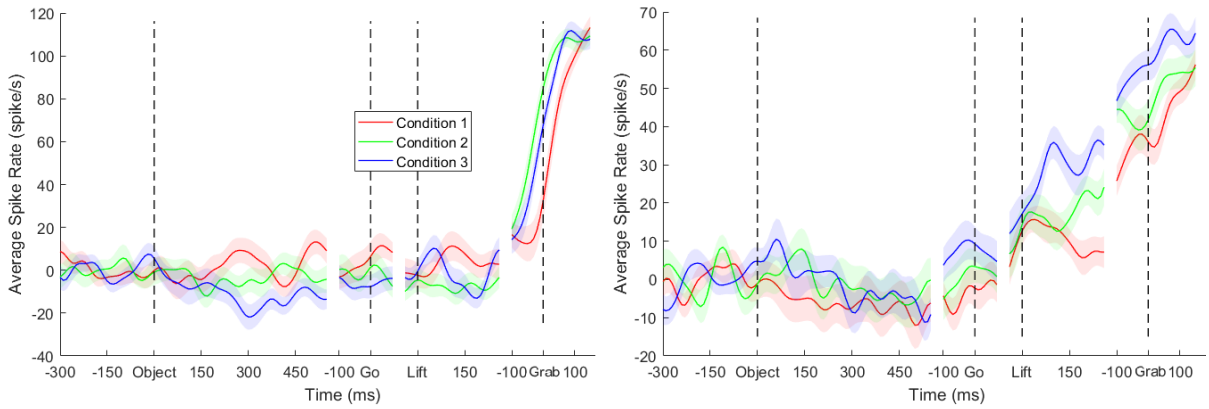

Figure 4.5: Filtered PSTH of the multi-unit activity detected in PMv of monkey J1 during a reach-to-grasp task: **(a)** channel 16 and **(b)** channel 43. The trials are aligned to four events of interest, and each trace plots the average spike rate across trials for one experimental condition. The standard error is plotted in a lighter shade around the PSTH for each condition.

4.2 Rotational structure in PMd and PMv dynamics

The spike rate of each channel can be used to create a 64-dimensional state-space, where each recorded channel represents one dimension. A possible state is represented by a point with 64 coordinates. An additional dimension - time - can be included to create the concept of trajectory, which is nothing more than a temporal sequence of states. As previously introduced in section 4.1.2, the neuronal activity is not independent for all channels; therefore, PCA is an

effective technique for dimensionality reduction. Still, most frequently, rotational patterns are not aligned to the axes of the state-space, and are, therefore, less obvious to the eye.

Following the work of Churchland *et al.* (2012), this section explores the dynamical structure in PMd and PMv neurons using jPCA. This method aims to find a new orthonormal basis that spans the same space as the input, but makes rotational structure more detectable. (8)

4.2.1 State-space modelling with jPCA

In the following analysis, fig. 4.6 and fig. 4.7 represent the rotational dynamical structure for both motor areas, PMd and PMv, for five different intervals of 300 milliseconds - starting at 300 ms before object onset, at 100 ms after object onset, at 50 ms after the go cue, at 150 ms before the monkey lifts the hand, and also at 200 ms before pulling the object -, as stated in table 4.1.

Table 4.1: Intervals used in the comparative analysis of the dynamical structure. The column ≪Alignment≫ defines the event to which the starting and ending points are relative to.

Interval	Starting point (ms)	Ending point (ms)	Alignment
I. Before object onset	-300	0	object onset
II. After object onset	100	400	object onset
III. After go cue	50	350	go cue
IV. During lift hand	-150	150	lift hand
V. During grab object	-200	100	pull object

The length of the intervals was chosen in order to display as much dynamical behaviour as possible, for a better comparison between areas and epochs without an overload of information in the plots. The limits of each interval were set so as to compare the rotational dynamics before and during the reach-to-grasp task (interval I versus intervals II-V), to observe the effect of a cue on the rotational structure (intervals II and III), and to monitor the evolution of the dynamical activity during important time points of the actual movement (intervals IV and V). As the neuronal response to a cue is not instantaneous, the analysis of their effect starts a few milliseconds after the cue itself.

Each line in a jPCA plane, like fig. 4.6 and fig. 4.7, represents the projection of a trajectory of neuronal states (from the starting point, identified with a circle, to the finishing instant, represented by an arrow) for a certain experimental condition. It might be tempting to view jPCA projections as spike rates, but that conception is wrong as these projections are found in a space that is not the high-dimensional space constituted by the recorded channels (which

can, indeed, be represented by spike rates). For this reason, jPCA projections come in arbitrary units (a.u.).

Comparing the activity between and within jPCA planes is certainly not trivial and requires an understanding of what exactly a rotation is and what selective planes look like. A trajectory is said to have strong rotational structure if the angle between itself and its derivative tends to $\pi/2$ radians – it may be intuitive to think of a circular trajectory, where the velocity vector is perpendicular to the displacement. Perhaps less trivial and harder to parameterise is the selectivity of a jPCA plane: a plane is more selective, the further the projections are in every instant, as this means distinct conditions are organised differently in state-space. It is also crucial to keep in mind that the projections are normalised within the interval; as such, it is not reasonable to compare the values of the projections between different epochs.

The following paragraphs are dedicated to describing the behaviour of PMd trajectories in state-space - *see* fig. 4.6.

Before receiving any stimulus, the dynamical behaviour in PMd (fig. 4.6(a)) is characterised by longer trajectories than in later epochs, which indicates that the neuronal state is changing more rapidly; also, these curves do not show a clear direction yet. This leads to strong overlaps of rotations, which make the conditions less distinguishable. The data variance captured in these three planes sums up to approximately 75%.

When the object becomes visible and the monkey starts planning the reach-to-grasp task, the neuronal activity assumes a more organised behaviour in the state-space (fig. 4.6(b)). The rotations are still strongly present, but the lines overlap much less, compared to the previous epoch. This can be interpreted as the neuronal state of different conditions spanning the state-space differently. It should, nevertheless, be intuitive that, at object onset, the three conditions are described by very similar neuronal states, as there is no motor planning occurring at that time. In fact, 100 ms after this cue, the projections of the state onto $jPCs_{1-3}$ are still relatively close, and progressively diverge over time. It is also interesting to note that, in jPCA plane 3, the projections do not coincide at any time instant and show locally strong rotations around distinct points. During this epoch, the model explains 87% of data variance, with a stronger emphasis on jPCA plane 3.

After the monkey is instructed to reach for the object, the projections of the neuronal state in PMd onto the jPCA planes (fig. 4.6(c)) are very similar to the previous epoch. In comparison to fig. 4.6(b), the projections in the first plane have a more accentuated rotational structure, which captures more data variance; and the projections in the third plane follow the opposite tendency. The most striking difference is that in this epoch the initial states (indicated with a circle) are very dissimilar and remain so throughout the whole interval. Overall, the variance

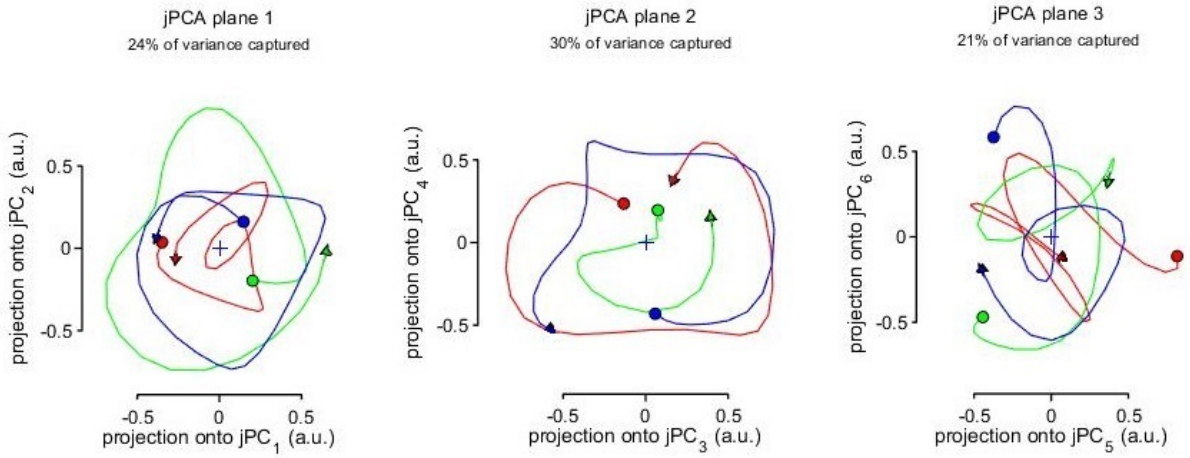

(a) before object onset

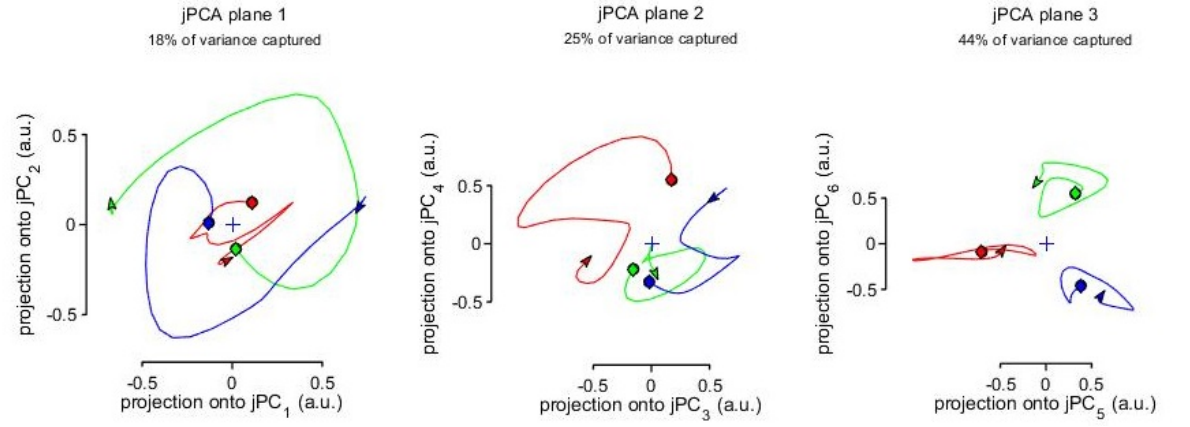

(b) after object onset

Figure 4.6: Projection onto state-space of the multi-unit activity of a PMd neuronal population in monkey J3. The 5 intervals analysed, **(a)**-**(e)**, are described in table 4.1. The 3 jPCA planes span the space defined by the 6 principal components that capture the most variance. The planes are ordered by decreasing value of the eigenvalues. Each trace represents the trajectory for condition 1 (in red), condition 2 (in green) and condition 3 (in blue). The circle and the arrow represent the starting and the end points of the analysis, respectively, and the diamond is plotted on the point corresponding to the event of interest.

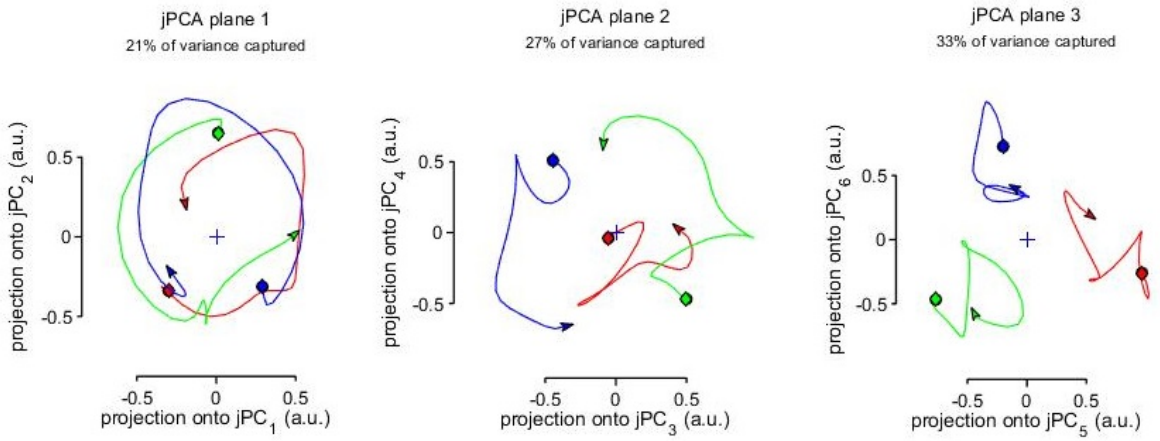

(c) after go cue

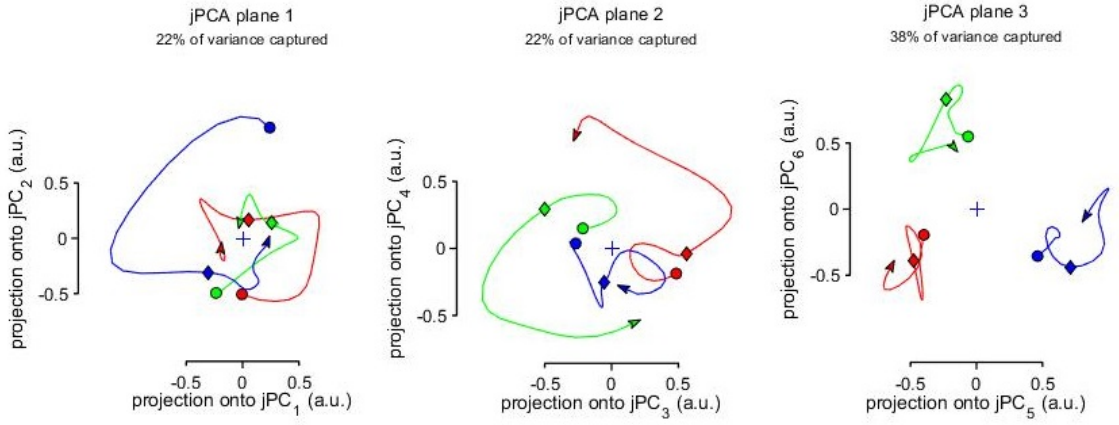

(d) during lift hand

Figure 4.6 (Cont.): Projection onto state-space of the multi-unit activity of a PMd neuronal population in monkey J3. The 5 intervals analysed, **(a)**-**(e)**, are described in table 4.1. The 3 jPCA planes span the space defined by the 6 principal components that capture the most variance. The planes are ordered by decreasing value of the eigenvalues. Each trace represents the trajectory for condition 1 (in red), condition 2 (in green) and condition 3 (in blue). The circle and the arrow represent the starting and the end points of the analysis, respectively, and the diamond is plotted on the point corresponding to the event of interest.

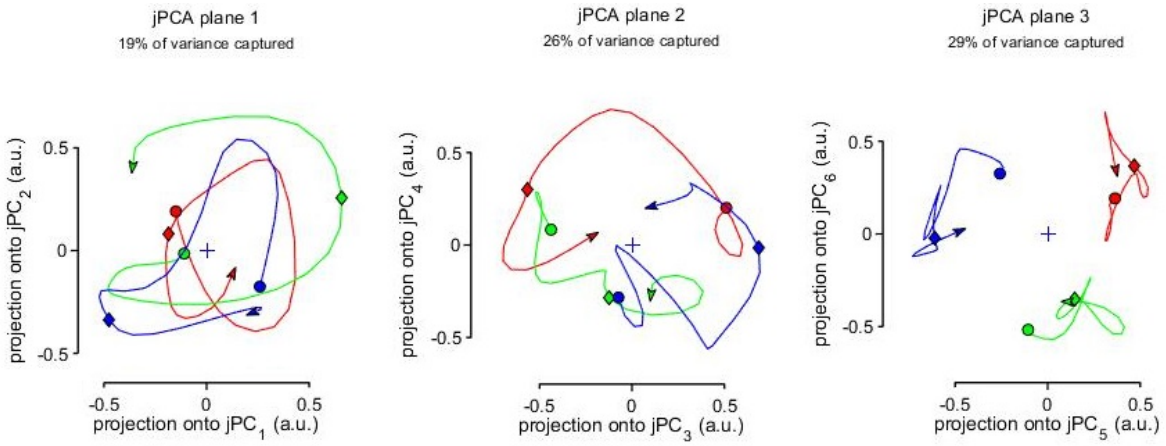

(e) during grab object

Figure 4.6 (Cont.): Projection onto state-space of the multi-unit activity of a PMd neuronal population in monkey J3. The 5 intervals analysed, **(a)**-**(e)**, are described in table 4.1. The 3 jPCA planes span the space defined by the 6 principal components that capture the most variance. The planes are ordered by decreasing value of the eigenvalues. Each trace represents the trajectory for condition 1 (in red), condition 2 (in green) and condition 3 (in blue). The circle and the arrow represent the starting and the end points of the analysis, respectively, and the diamond is plotted on the point corresponding to the event of interest.

captured by the three jPCA planes for the present epoch is 81%.

The trajectories of the neuronal state in PMd during the actual lift of the hand (fig. 4.6(d)) include activity before and after the movement starts, which makes it particularly interesting to inspect. In the first jPCA plane, the projections after the lift of the hand (between the diamond and the arrow) are notably more clustered than before the movement (between the circle and the diamond). Conversely, for the second and third planes, the states are noticeably more clustered before the movement starts. These projections capture 82% of the overall data variance, with emphasis on the third plane.

The state-space representation of the neuronal activity in PMd during the pull of the object (fig. 4.6(e)) seems to have partially lost its organisation, comparing to the previous epoch, resembling the behaviour observed before the object onset. In fact, the first two planes display less distinct trajectories, and the third plane is mostly characterised by straight lines. It is worth mentioning that, after the monkey grabs the object (illustrated by a diamond), the neuronal states evolve inwards, probably bringing the neuronal states together, back to the baseline activity. Altogether, these projections capture 74% of the data variance.

A similar description can now be made for the dynamic trajectories of PMv's neurons in state-space - *see* fig. 4.7.

Similar to PMd, the neuronal state in PMv (fig. 4.7(a)) before the monkey sees the object is

very disorganised, with very long trajectories, which indicates that two consecutive states are very different from each other during this interval. The projections of different conditions do not show distinct patterns; on the contrary, they are highly overlapped. These three planes capture 78% of data variance evenly.

After signaling the position of the object, the dynamical behaviour of the neurons in PMv (fig. 4.7(b)) starts to acquire a more condition-dependent structure, which is especially notice-

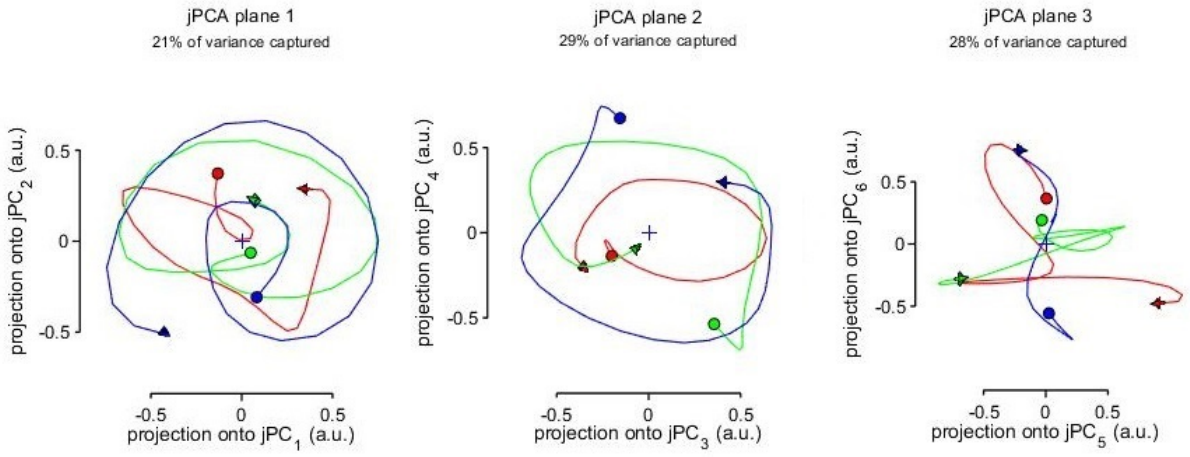

(a) before object onset

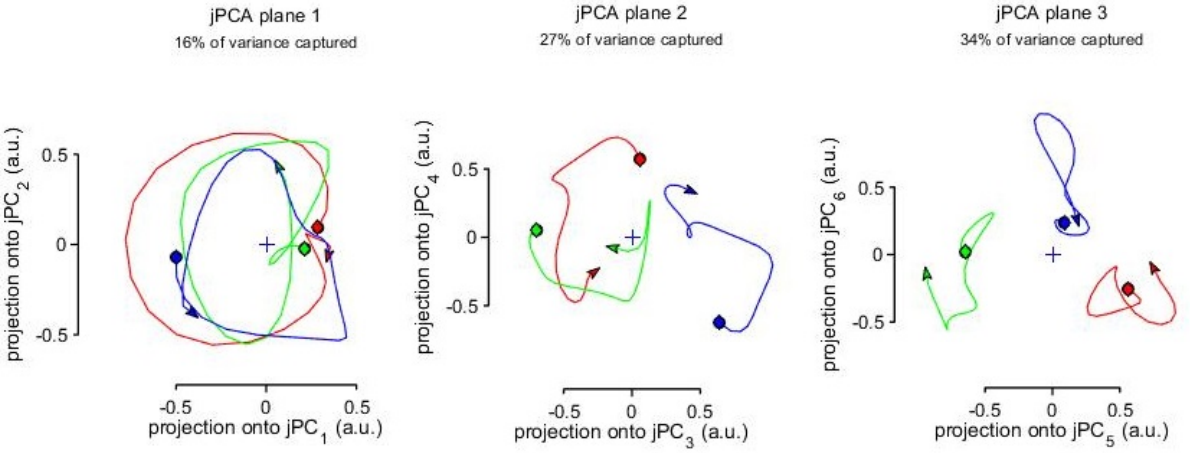

(b) after object onset

Figure 4.7: Projection onto state-space of the multi-unit activity of a PMv neuronal population in monkey J3. The 5 intervals analysed, **(a)**-**(e)**, are described in table 4.1. The 3 jPCA planes span the space defined by the 6 principal components that capture the most variance. The planes are ordered by decreasing value of the eigenvalues. Each trace represents the trajectory for condition 1 (in red), condition 2 (in green) and condition 3 (in blue). The circle and the arrow represent the starting and the end points of the analysis, respectively, and the diamond is plotted on the point corresponding to the event of interest.

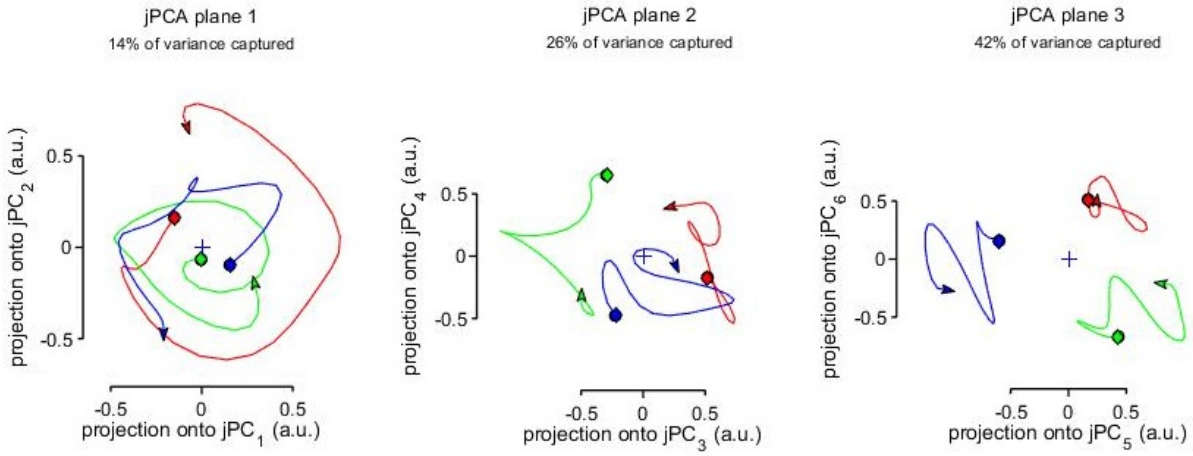

(c) after go cue

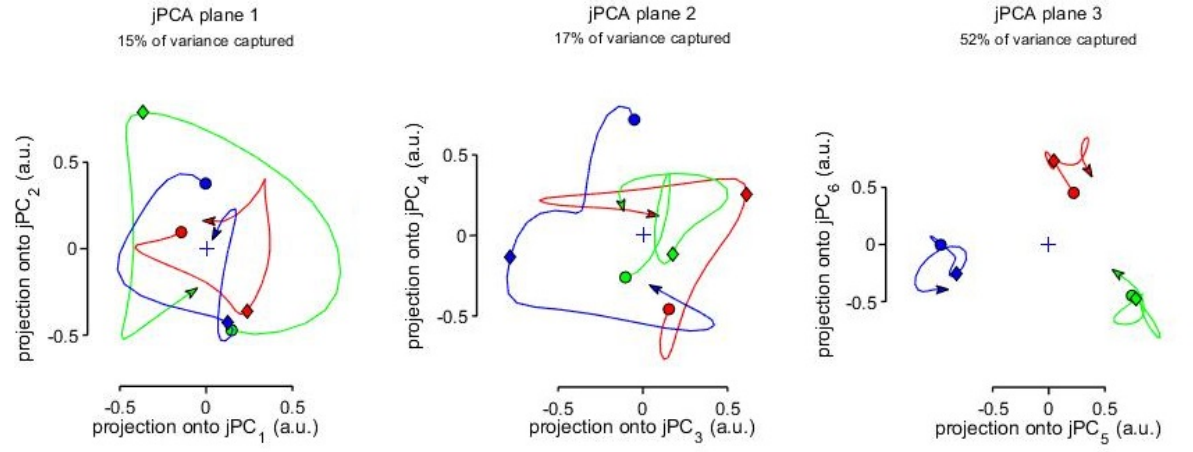

(d) during lift hand

Figure 4.7 (Cont.): Projection onto state-space of the multi-unit activity of a PMv neuronal population in monkey J3. The 5 intervals analysed, **(a)**-**(e)**, are described in table 4.1. The 3 jPCA planes span the space defined by the 6 principal components that capture the most variance. The planes are ordered by decreasing value of the eigenvalues. Each trace represents the trajectory for condition 1 (in red), condition 2 (in green) and condition 3 (in blue). The circle and the arrow represent the starting and the end points of the analysis, respectively, and the diamond is plotted on the point corresponding to the event of interest.

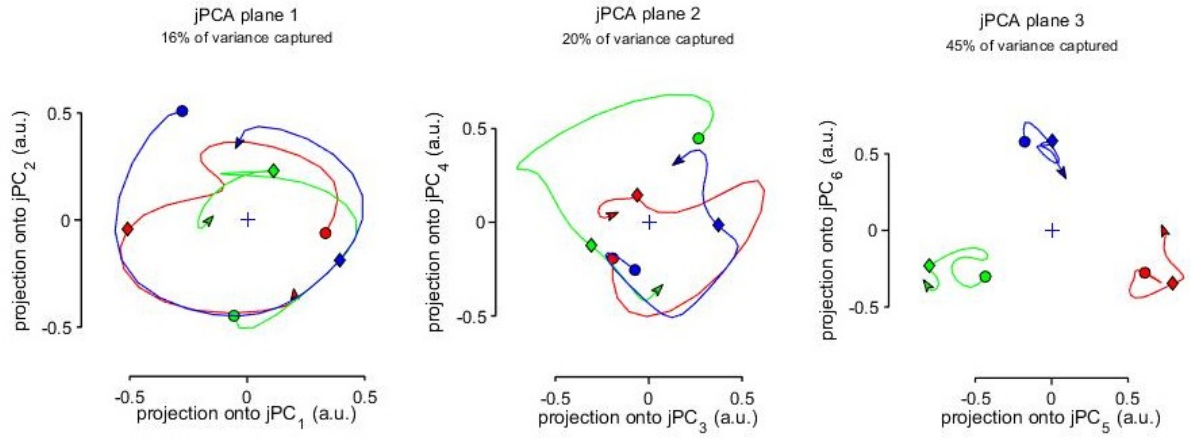

(e) during grab object

Figure 4.7 (Cont.): Projection onto state-space of the multi-unit activity of a PMv neuronal population in monkey J3. The 5 intervals analysed, **(a)**-**(e)**, are described in table 4.1. The 3 jPCA planes span the space defined by the 6 principal components that capture the most variance. The planes are ordered by decreasing value of the eigenvalues. Each trace represents the trajectory for condition 1 (in red), condition 2 (in green) and condition 3 (in blue). The circle and the arrow represent the starting and the end points of the analysis, respectively, and the diamond is plotted on the point corresponding to the event of interest.

able on the jPCA planes 2 and 3, where the projections do not overlap as much. Following the same reasoning as for PMd, one would expect to see clustered initial states; these might not be evident in all projections, but, indeed, the initial states are very similar in jPC_2 and jPC_6. The variance captured by these three planes sums up to 77%, with emphasis in the last two jPCA planes.

Regarding the projections of the neuronal state in PMv after the go cue (fig. 4.7(c)), all three planes have strong rotations and there is a clear distinction in the pattern of activity for each condition, with very few overlaps in the projections. It is also visible that the initial state is very similar for all conditions in the first jPCA plane. Additionally, the distribution of the variance explained by each plane is very uneven, with the third jPCA plane capturing most of the variance, which sums up to a total of 82%.

During the lift of the hand, PMv's trajectories in state-space (fig. 4.7(d)) are extremely disparate between planes. Although rotational structure and distinct patterns of activity are present, the first two jPCA planes capture very little variance of the data. The third plane captures almost two thirds of the 84% of the data variance explained in these planes. By the end of the epoch, the projections are oriented inwards, almost as if the neuronal state was heading towards the baseline activity.

In the next epoch, however, the rotational structure of the projections ((fig. 4.7(e)) is still

robust, and the patterns of activity remain very distinct between conditions. This structure contrasts with the same epoch for PMd, where rotational structure weakens after the grasp, but is consistent with the PSTH plot (fig. 4.5), which shows strong condition-selectivity during this interval. The jPCA planes capture a total of 81%, once more mostly explained by the third plane.

Despite not being an adequate indicator of goodness of fit of the jPCA method to data, it is a good a idea to make sure that a high fraction of the data variance is represented in the low-dimensional space. This decreases the likelihood of analysing dynamical structure that is less representative of the neuronal activity. Indeed, as depicted in fig. 4.8, the 6-dimensional space found captures around 80% of the data variance for all the datasets and intervals studied.

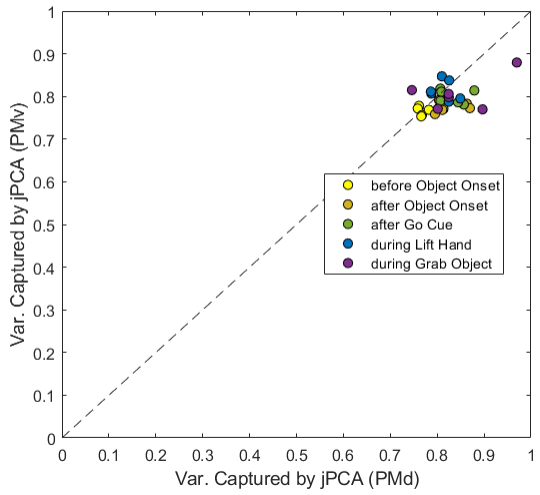

Figure 4.8: Comparison between PMd and PMv of the variance captured by the 6 jPCA projections for all datasets in all intervals mentioned in table 4.1. The model was calculated with the rank-6 matrices that capture dynamics in all 6 analysed dimensions.

4.2.2 Goodness of fit of the jPCA model

To assess the quality of the fit of the constrained model, M_{skew} to both PMd and PMv, the coefficient of determination (CD), R^2, can be used to express how well rotational dynamics predict the overall data behaviour. Additionally, rotational goodness-of-fit ratio (RGR) can also be of useful to understand the fraction of linear dynamics that consist on rotations.

Figure 4.9 compares the fit quality of the model for PMd and PMv. This analysis is performed for all the six datasets available and all the 5 epochs mentioned in table 4.1.

At first sight, the distribution of the points in fig. 4.9(a) does not show evidence that the model fits the activity of one area better than the other. However, when looking at each interval individually, there seems to be a very slight tendency towards one of the sides of the diagonal

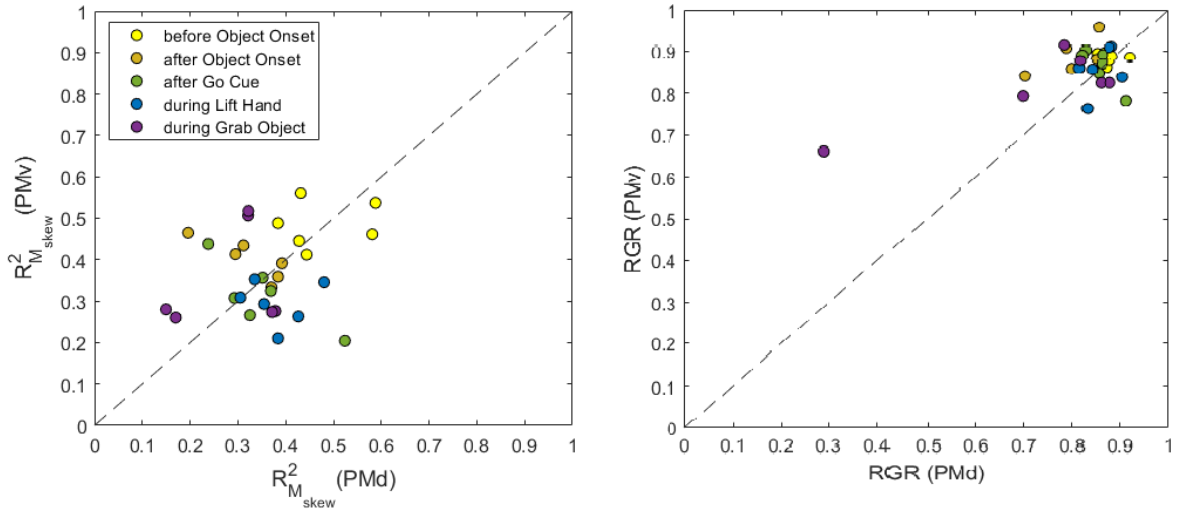

(a) **(b)**

Figure 4.9: **(a)** Comparison of the coefficient of determination of the constrained systems in PMd and PMv. This plot represents datasets J1-6 for the time intervals in table 4.1. **(b)** Comparison of the coefficient of determination ratio between the constrained, M_{skew}, and the unconstrained, M, systems, $R^2_{M_{skew}}/R^2_M$, in PMd and PMv. This ratio indicates how much of the total linear dynamics are explained by rotations in state-space. The coefficient of determination was calculated for the rank-6 matrices that capture dynamics in all 6 analysed dimensions.

line, suggesting that, for different intervals, the model might fit the activity in one premotor area better than the other.

Similarly, the RGR between the unconstrained and constrained models in fig. 4.9(b) does not show considerable difference between the two motor areas. It is mostly evident from this plot that most of the variance explained by linear dynamics are captured by the rotations in the state-space.

4.2.3 Control analyses

A potential concern with the use of jPCA is that it may be sufficiently powerful to find rotational dynamics for any population with diversified and multiphasic responses. Such effect is particularly critical when analysing few conditions in a high-dimensional space. (8)

The present work analyses 3 conditions in a 64-dimensional space. As such, it is sensible to analytically assess the extent to which rotational patterns may be found by chance, when responses are diverse and complex but do not correspond to actual neuronal dynamics. A first evaluation, based on the analysis of Churchland et al. (2012), involves three shuffled controls that disrupt the deep structure of the data but conserve the covariance, as in, maintains the diversity and complexity of the responses. This shuffling test is particularly important when

analysing few conditions, as these tend to have less variability, making it more likely for random variations to be interpreted as important patterns. (8) The second test, inspired by the work of Lara et al. (2018), alters the covariance of the data and, consequently, the dynamical structure as well. (56)

4.2.3.A Disrupting the dynamical structure

Based on the knowledge that both PMd and PMv are involved in the planning of an action (4), these shuffled controls are built around the distinction between pre-stimulus phase and the stimulus-driven phase. For this purpose, the division between both periods is set at a time-point 100 ms after the object onset. At this instant, there should already be some condition-dependent activity, which will be disrupted by the shuffling. (8)

In the first shuffled control, the stimulus-driven activity was time-inverted for 1 or 2 conditions randomly selected, preserving the continuity of the signals for each neuron. For the second shuffled control, the pattern of stimulus-driven activity was time-inverted for all conditions, preserving the continuity of the signals for each neuron. The third shuffled control consists in permuting the stimulus-driven activity of different conditions. The stimulus-driven activity of each condition is necessarily appended to the final firing rate of the pre-stimulus activity of a different condition, preserving the continuity of the signals for each neuron. These manipulations are not expected to remove all rotational structure, but rather to decrease the consistency of any true rotations. (8) The effect of the shuffled controls on the PSTHs and the rotational structure of PMd and PMv recordings can be found in appendix B.

It is important to clarify one crucial concept: the rotational structure is said to be stronger or more robust the closer the angle between the trajectory and its derivative is to $+\,\pi/2$. If strong and robust rotational structure survives the shuffling operations, that should be an indication that the jPCA method can extract such structure when it is weakly present. Conversely, if the shuffle procedures cause rotational structure to be lost, that should validate the method, showing that this structure was present in the original data to a greater extent than it would by chance. (8)

To efficiently quantify rotation strength, one can measure the angle from the neural state on a jPCA plane, X, to its derivative, $\dot{X}$. For each jPCA plane, this angle is calculated for every condition and time point, and a histogram of the angles is built. The rotation robustness of the neural state is, thus, calculated as the sum of the three histograms of the neural state weighted by the variance captured by each plane. The resulting histogram is, then, normalised with the total variance captured by the planes.

Since the shuffled controls only disrupt the structure of the data from the point-instant at

100 ms after object onset on, the analysis of rotational strength in fig. 4.10 is performed for the interval $[90; 290]$ after object onset. This figure compares the rotational strength of the shuffled data and the original data.

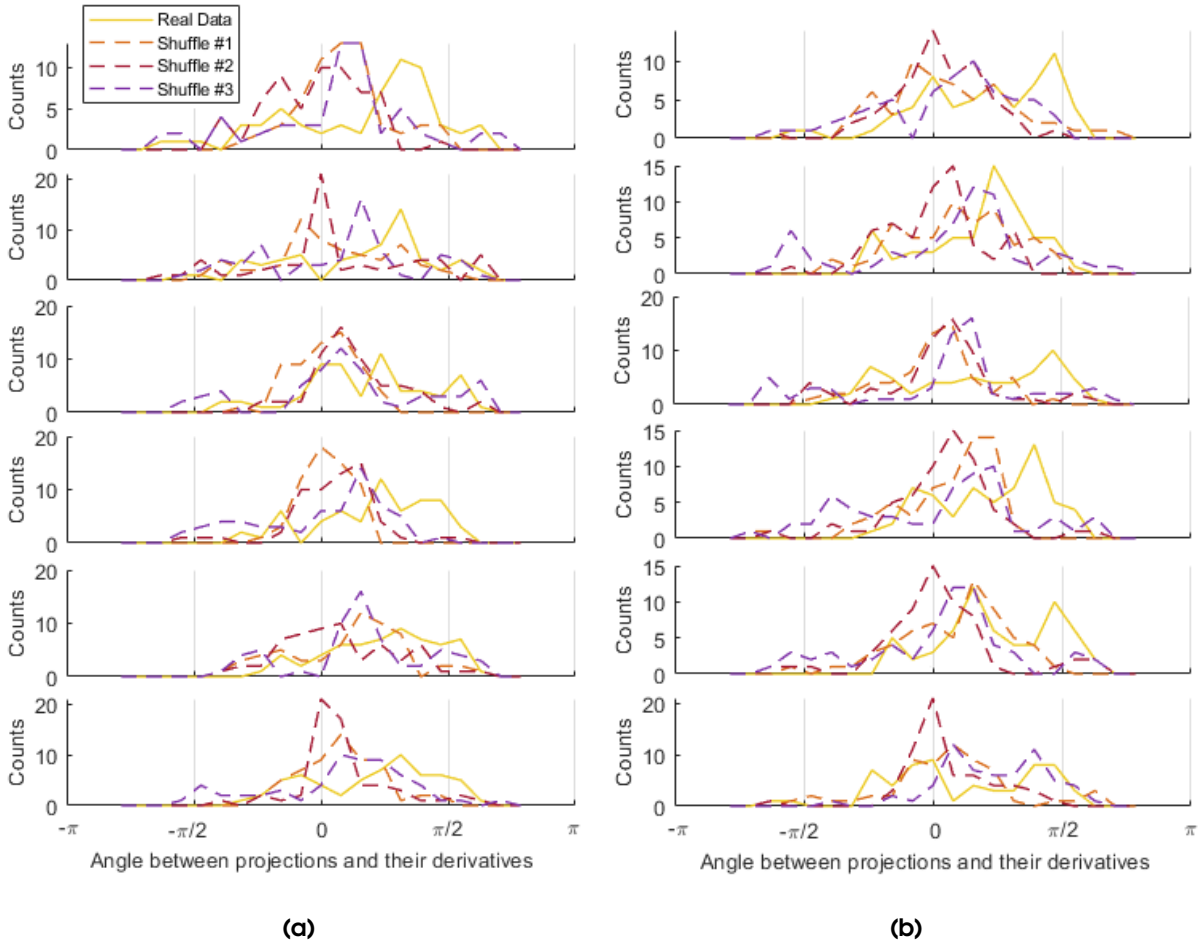

Figure 4.10: **(a)** Normalised weighted sum of the histograms of the angle between the neural state, X, and its derivative, $\dot{X}$, projected onto each of the 3 jPCA planes that capture the most variance. The weights for each histogram correspond to the variance captured by the corresponding jPCA plane. Each histogram corresponds to the angles observed in state-space in the interval (90; 290) after object onset for the original and shuffled data, for PMd neural populations in datasets J1-6, respectively from top to bottom. **(b)** Same as (a) for PMv neural populations.

Because the jPCA planes are oriented such that the rotational structure, however weak or strong, is projected anticlockwise, robust rotations should have positive angles with values close to $\pi/2$. Following the same reasoning, pure scaling and expansion in the trajectories result in angles near 0 and π. (8)

For most datasets, the angles of the trajectories in fig. 4.10(a) and fig. 4.10(b), respectively PMd and PMv, show a tendency for the original data to have trajectories with angles nearer

$\pi/2$, and weaker rotations for the shuffled controls. In some histograms, there are distinct peaks for the original and simulated data that illustrate more clearly that shuffled data have more scaling and expansion than the original data. Moreover, all original datasets show a complete absence of $-\pi/2$ angles, which is not true for the simulated data.

It is also predictable that, since the deep structure of the shuffled data is disrupted, the model loses some of its ability to predict $\dot{X}$ from X, and, as a consequence, projections lose some of its rotational structure to the detriment of gaining more linear structure. This, once again, can be measured recurring to the CD, $R^2_{M_{skew}}$, and the RGR, $R^2_{M_{skew}}/R^2_M$. Figure 4.11 plots the fit quality of the shuffled data and the original data for all datasets and for both areas, using both statistical measures.

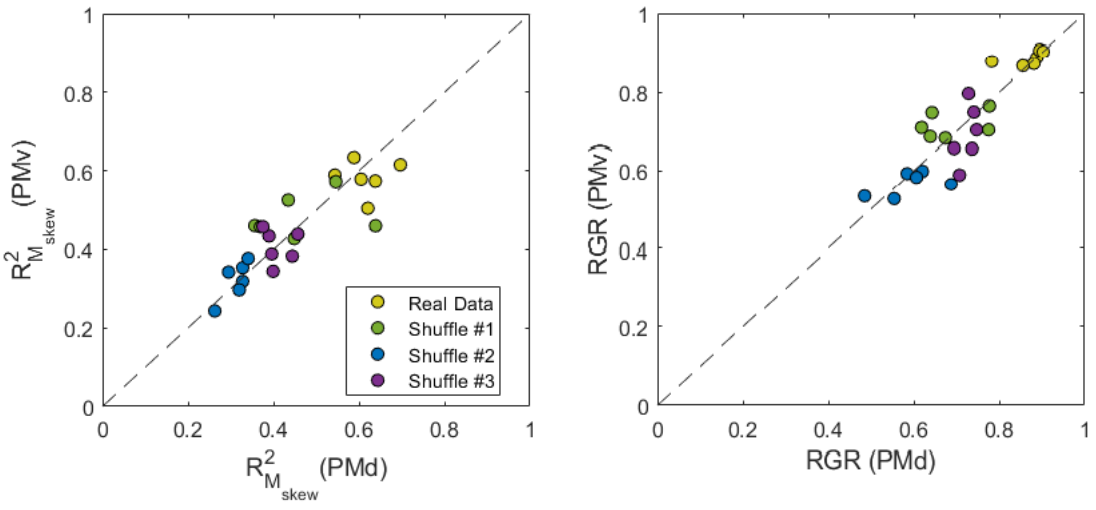

(a) **(b)**

Figure 4.11: **(a)** Comparison of the effect of the shuffled controls on the CD of the constrained system, $R^2_{M_{skew}}$, between PMd and PMv. **(b)** Similar to **(a)**, for the RGR between the constrained and unconstrained systems of two, $R^2_{M_{skew}}/R^2_M$. These analyses were performed for 200 ms after object onset. The coefficient of determination was calculated for the rank-6 matrices that capture dynamics in all 6 analysed dimensions. The circles plot performance for individual datasets J1-6 during 200 ms after Object Onset.

According to fig. 4.11(a), the CD of the jPCA model is undoubtedly better for the original data. Among the shuffled controls, the shuffle #1 is the one that fits the model the best, followed by shuffle #3. Regarding the relative quantity of rotational structure in the dynamical model in fig. 4.11(b), the original data surpasses the simulated data. Moreover, the data generated by the shuffled controls #1 and #3 have significantly more rotations that the trajectories in shuffle #2.

4.2.3.B Altering the data covariance

The previous control test accounts for the possibility that jPCA might be powerful enough to find rotational structure in state-space for populations with multiphasic responses. (8) However, it does not consider that "random" variations in trajectories of individual conditions might be the reason why jPCA fits the neuronal activity in these datasets so well. This is particularly important because the quality of the dynamical fit is largely determined by whether the dynamical trajectories of distinct conditions obey similar flow fields. (56)

To test that, a bootstrap is employed in which both the dynamical structure and the data covariance are disrupted. Each of the 6 columns of $\dot{X}$ and X is modified to include 3 redrawn conditions, with replacement, from 3 columns, also with replacement. To avoid two identical columns, two entries in the same (modified) row cannot contain data from the same (original) column. This process is repeated 1000 times for each dataset and epoch to provide a sampling distribution. The *p*-values for any comparison is the percentage of draws for which the statistical measure, $R^2_{M_{skew}}$ or $R^2_{M_{skew}}/R^2_M$, is greater for the modified data than the original data.

For most datasets and intervals, the *p*-values in both plots in fig. 4.12 were lower than 0.05%. A low *p*-value for the analysis with $R^2_{M_{skew}}$ signalises a disruption in the dynamical structure of

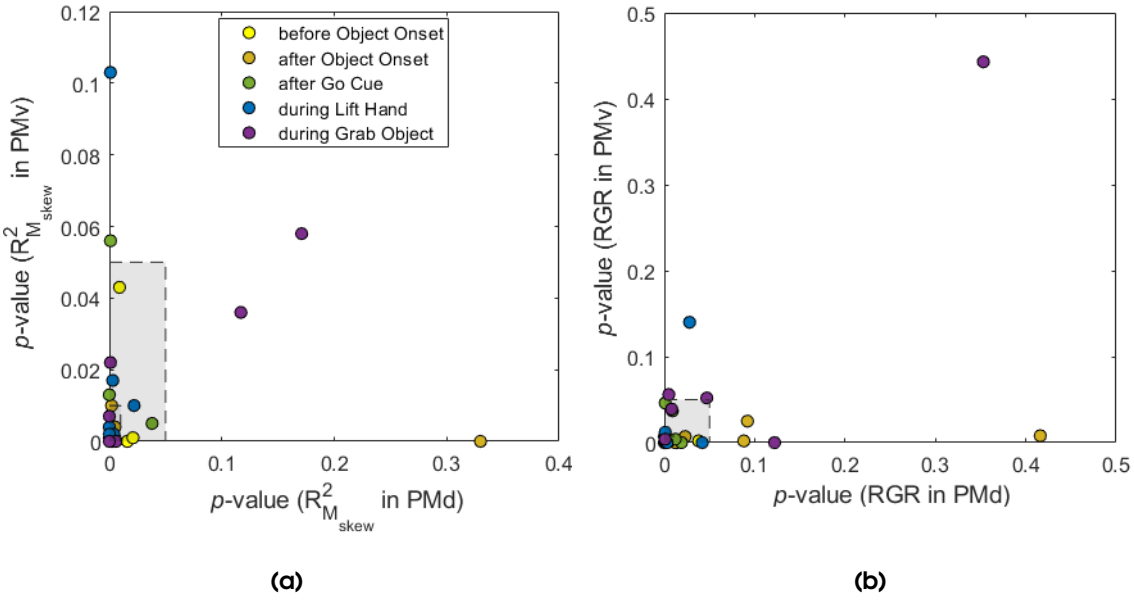

Figure 4.12: **(a)** Comparison of the statistical significance $R^2_{M_{skew}}$ between PMd and PMv. The *p-values represent the likelihood of finding a random population in 1000 bootstraps with a higher value than that of the original recordings. Each bootstrap involved modifying each of the 6 columns of X and $\dot{X}$ to include data from 3 redrawn condition with replacement. This analysis was performed for the same intervals referred to in table 4.1. The shaded area represents the region where p-value <0.05.* ***(b)*** *Similar to (a), for the RGR, $R^2_{M_{skew}}/R^2_M$. The model was calculated with the rank-6 matrices that capture dynamics in all 6 analysed dimensions.*

the data that led to either a loss of rotational structure or a big dissimilarity in flow field of each condition. A low *p*-value in the analysis of the RGR, $R^2_{M_{skew}}/R^2_M$, should be interpreted as a loss of rotational structure relative to the general linear dynamics.

Notably, there are some points that display significant results for one area but not for the other, and one single point representing a non-significant (*p*-value above 0.05) result for both areas. Although not displayed in fig. 4.12, the four points most distant from the shaded area (1 blue, 2 purple and 1 yellow) correspond to the same datasets and intervals in both plots. These points likely represent outlying activity.

4.3 Unraveling the neuronal code in PMd

PMd populations have previously been described to contain rotations in their state during non-periodic movements, like the reach-to-grasp task. This seemingly simple structure can explain complex features of movement. (8)

In a simplistic fashion, the control of arm movement via a BMI requires the clear identification of the movement direction and timing of the action. In this context, this section attempts to give some intelligibility to the rotational structure in PMd.

It is relevant to clarify that the initial intent was to perform the following analysis for PMv as well, but technical difficulties precluded the data acquisition for both areas. As an alternative, previous PMd recordings are going to be used.

4.3.1 Patterns in neuronal activity

Following the context principle of neuronal activity, neurons modulate their activity according to the event. Moreover, the degeneracy principle proposes that many neurons can encode the same behaviours.

The principles are demonstrated in fig. 4.13. This figure illustrates four examples of neurons, whose individual activity is tuned to the trial characteristics. Each of the four neurons shows similar patterns for the trials recorded under the same conditions, and some level of discrepancy relative to other conditions, at least at some point after object onset. Furthermore, fig. 4.13(d), in particular, tells apart the visual and the motor phases.

Although not all recorded neurons show such clear condition- and epoch-tuning, many of them do. It is, however, not clear what this implicates in terms of projection onto the jPCA planes. But, following the same reasoning, one can hypothesise is that rotational structure will aggregate for the trials recorded under the same conditions.

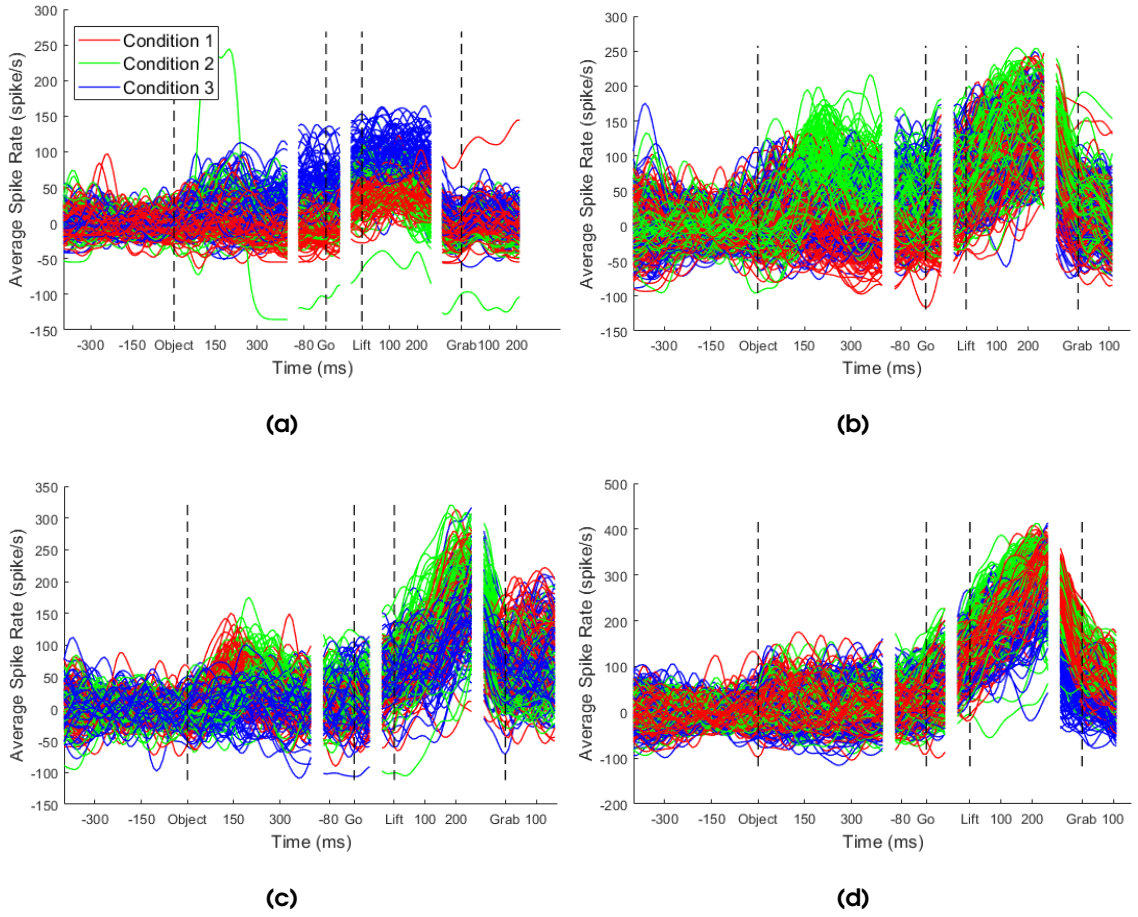

Figure 4.13: Neuronal activity recorded for all trials in a given dataset for four different channels: **(a)** channel 79 of dataset K1, **(b)** channel 37 of dataset K5, **(c)** channel 85 of dataset K4 and **(d)** channel 90 of dataset K6. The trials are aligned to four events of interest, and each trace plots one trial. Trials recorded under different conditions are represented with different colours.

To test this, the approach to the jPCA method has to be slightly modified. It is now necessary to input each trial individually, rather than the across-trial average of the trials with the same condition, as in section 4.2. In principle, this should not make a big difference in the jPCA planes found, as the variation between trials and the across-trial average should consist in noise and artefacts, which should, nevertheless, be disregarded after PCA. For this reason, the trials recorded under the same conditions should have a similar flow field in the 6-dimensional space.

Rotational behaviour was analysed for the three intervals presented in table 4.2: one control interval with no activity triggered by any stimulus, one interval after the object onset, and another one starting shortly before the beginning of the movement.

In fact, as represented in fig. 4.14 before the subject visualises the object, the rotational structure is not organised, and there is no sign of a possible pattern. However, right after the

Table 4.2: Intervals used in the decoding of movement direction. The column ≪Alignment≫ defines the event to which the starting and ending points are relative to.

Interval	Starting point (ms)	Ending point (ms)	Alignment
I. Before object onset	-300	0	object onset
II. After object onset	50	350	object onset
III. During lift hand	-50	250	lift hand

object onset, recognisable patterns can be detected, as trials with the same conditions seem to have a similar spatial organisation especially in jPCA planes 2 and 3. This arrangement is even more prominent later in the task, during the actual lift of the hand. It is worth noting that the connection between strong rotations and trial condition is clearer during the actual movement, when the stronger rotations of jPCA planes 1 and 2 display a more distinct organisation.

The question now is whether the differences in rotational structure are enough to collect interesting information in the context of motor control for BMIs.

4.3.2 Decoding movement direction

This section is intended to evaluate if distinct movement directions - i.e., recording condition - feature different rotational structure.

4.3.2.A Input to Support Vector Machine (SVM)

Given the previous results, it might be fruitful to attempt the decoding for each of the intervals in table 4.2.

The first interval serves as a negative control, as no distinction is expected between the neuronal activity associated to each condition before object onset. This is true for raw spike rates and for the jPCA projections. An example of the jPCA-modelled input is shown in fig. 4.15(a).

As previously showed, jPCA projections do not explain the entirety of neuronal activity variance, making it interesting to decode the same datasets during the same periods but with the raw neuronal activity of all channels, as exemplified in fig. 4.15(c). This feature design can be compared with the dynamical approach defined by the jPCA method for purpose of understanding if, indeed, a low-dimensional and dynamical representation of neuronal activity can describe motor parameters better than the MUA itself.

Most importantly, fig. 4.15(b) illustrates an example of the jPCA-modelled input to the SVM, the true focus of this analysis.

For each of the input observations, each feature corresponds to the value of the projection

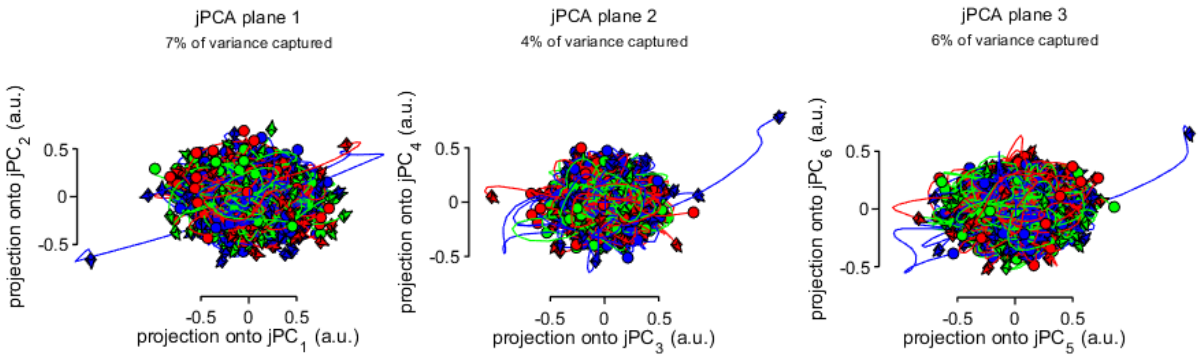

(a) before object onset

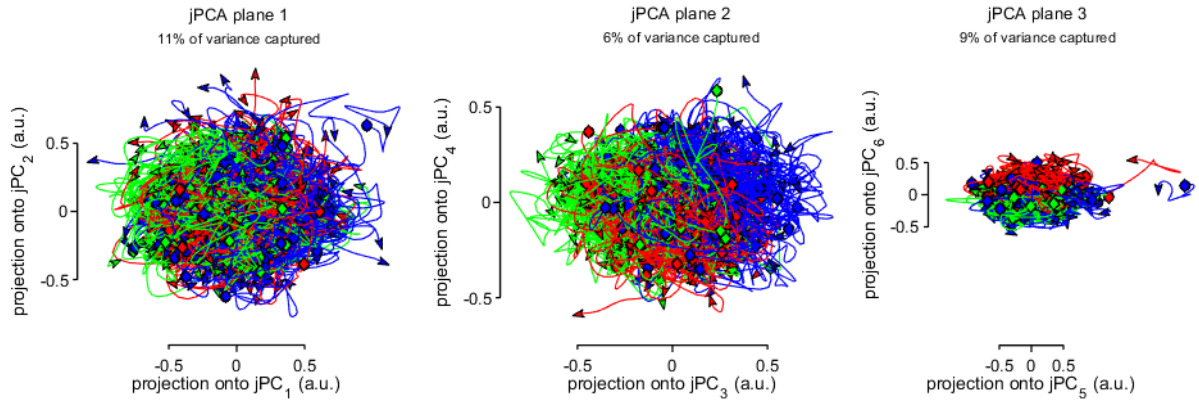

(b) after object onset

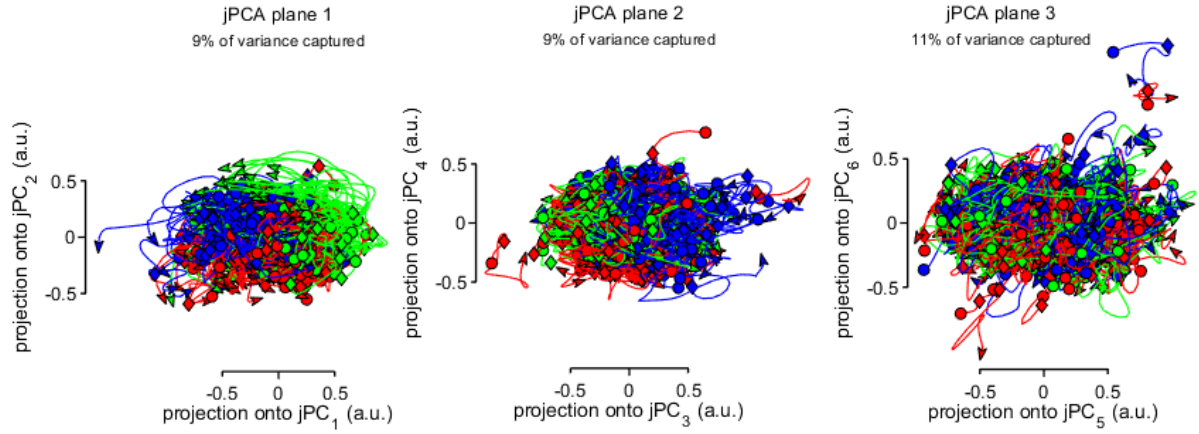

(c) during lift hand

Figure 4.14: Projection onto state-space all the trials recorded in monkey K8, referring to the multi-unit activity of a PMd population. The intervals analysed are: **(a)** (-300; 0) aligned to object onset, **(b)** (50; 350) aligned to object onset and **(c)** (-50; 250) aligned to lift hand. The 3 jPCA planes span the space defined by the 6 principal components that capture the most variance. The planes are ordered by decreasing value of the eigenvalues. Each trace represents the trajectory for a single trial, associated with condition 1 (in red), condition 2 (in green) or condition 3 (in blue). The circle and the arrow represent the starting and the end points of the analysis, respectively.

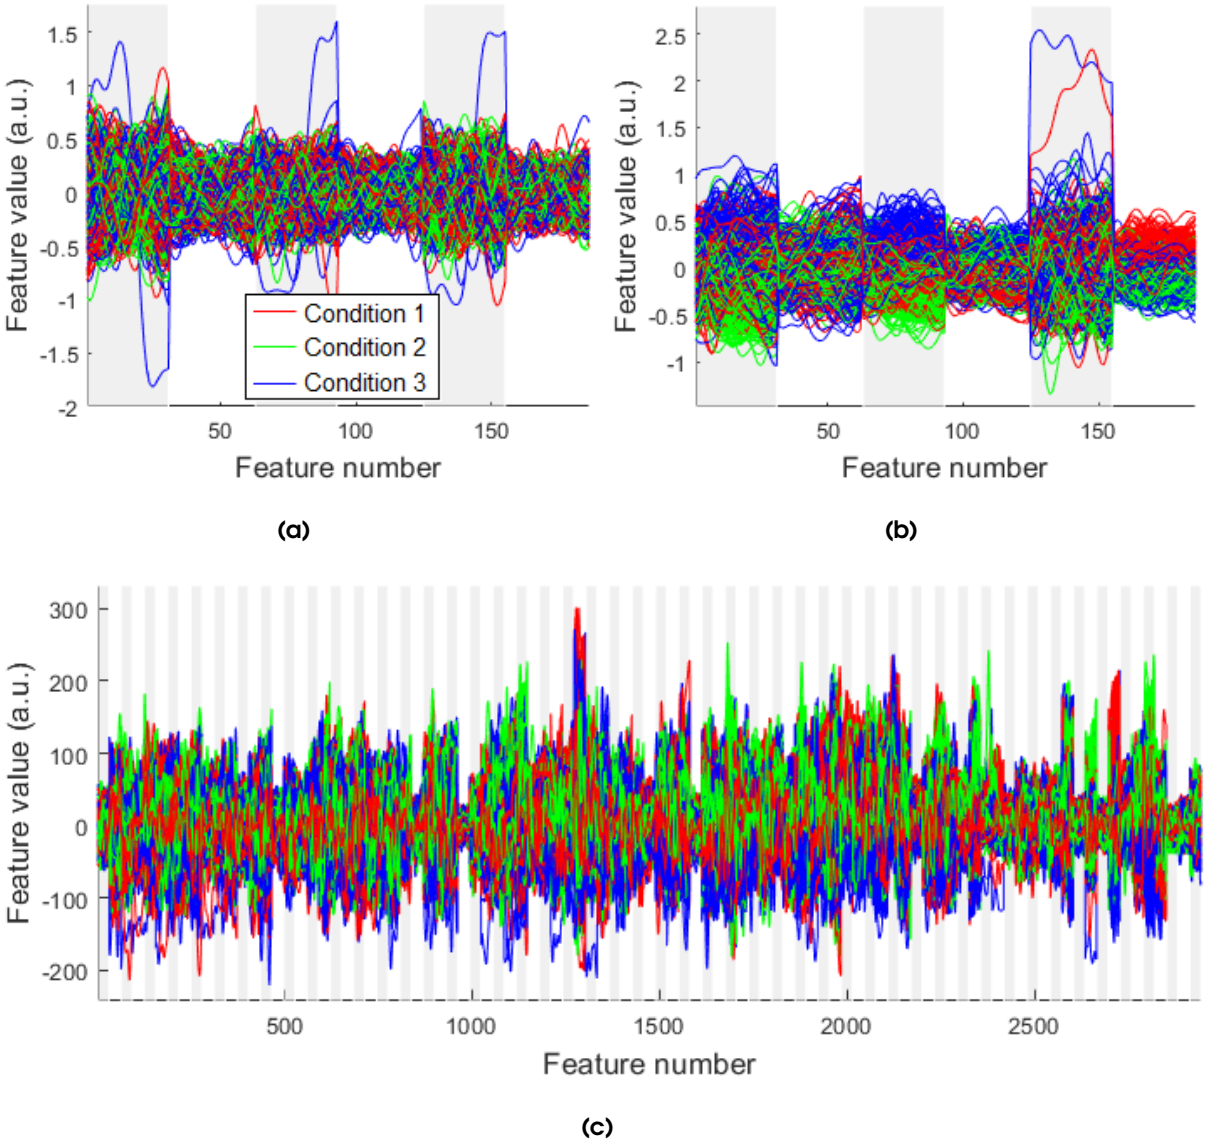

Figure 4.15: Example of the input observations used to decode and control the decoding of movement direction in monkey K8. **(a)** Features are the amplitude of the jPCA projections before object onset (interval I) at each instant - negative control. **(b)** Features are the amplitude of the jPCA projections after object onset (interval II) at each instant. **(c)** Features are the rates of every recorded channel after object onset (interval II) at each instant. The grey-shaded areas identify the odd jPCA projections (for **(a)** and **(b)**) and channels (for **(c)**).

(or raw spike rate, for the analysis with the entire neuronal ensemble) at a different time instant. Each of the intercalated grey- and white-shaded areas in the plots of fig. 4.15 correspond to a different jPCA projection (or recorded channel).

4.3.2.B Decoding accuracy

The distribution of the accuracy of decoding movement direction is represented in fig. 4.16 for each dataset, interval, and feature design (jPCA or the entire neuronal ensemble).

UPDATE FIGURES WITH LAST DATASET!

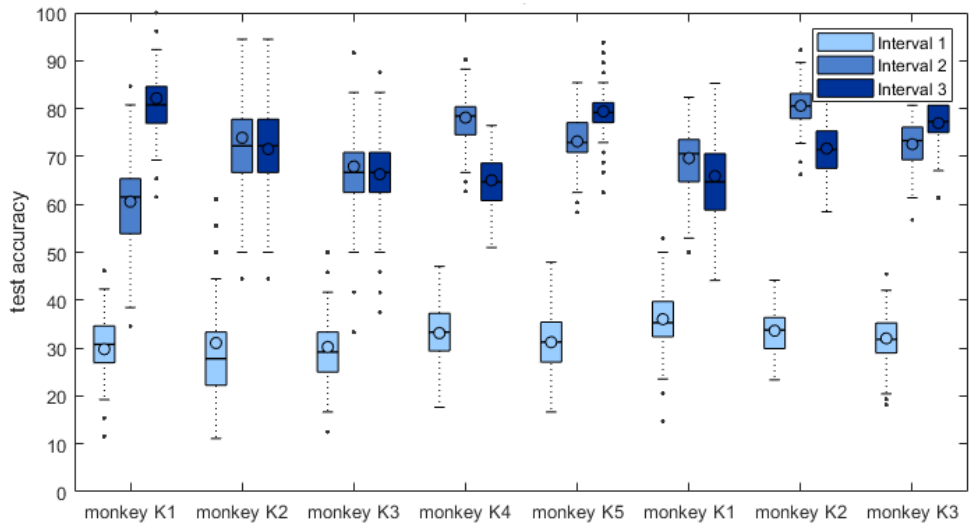

(a)

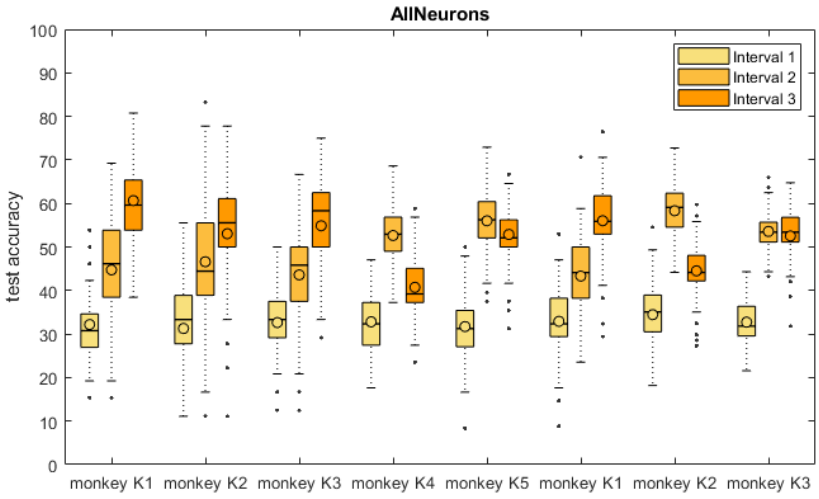

(b)

Figure 4.16: Accuracy of decoding movement direction with **(a)** jPCA projections versus **(b)** the entire neuronal ensemble. The decoding procedure was performed for datasets K1-9 for the time intervals in table 4.2. Each combination of dataset and interval was decoded 100 times to obtain the distribution. The boxes encompass the values between the 25th and 75 percentiles. The line and circle inside the box represent the median and mean values of the distribution, respectively. The fences limit the minimum and maximum values not considered outliers.

To obtain sufficient results to plot the accuracy distribution, the SVM was trained 100 times for each combination of dataset, interval and feature design. In each iteration, the machine is trained with a different set of 90% of the total observations in the datasets, and validated with the remaining trials.

The experimental setup is designed so that each object (experimental condition) is cued as often as all the other objects. Thus, the chance level for randomly attributing the correct label to an observation in the test set is about 33%. This value is not rigorous, as each test set is not constrained to have a balanced number of observation with each label.

Indeed, the mean and median accuracy of the decoding with either feature design, in figs. 4.16(a) and 4.16(b) was centred around 33% for interval 1 (before object onset). For the other two intervals, the accuracy was significantly above the chance level. Some datasets were more informative for one interval than for the other, which, curiously, was consistent between the two feature designs.

The most interesting result is that the decoding performance was undoubtedly superior for the feature design with jPCA projections than with the all the individual channels.

UPDATE FIGURES WITH LAST DATASET!

Thinking about a possible application to real-time BMI control, fig. 4.17 plots the training time. This interval corresponds to the time necessary for the SVM to identify the most important

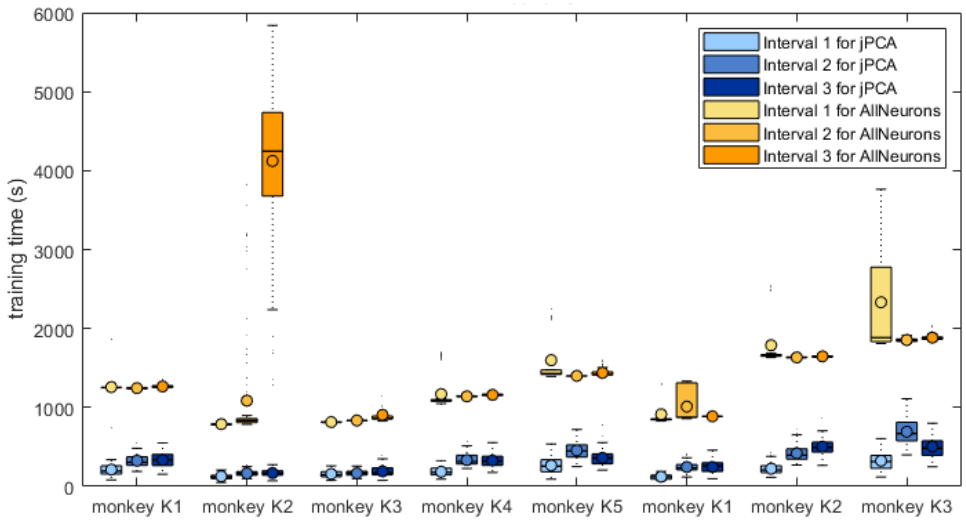

Figure 4.17: Time necessary to train the SVM for decoding movement direction with **(a)** jPCA projections versus **(b)** the entire neuronal ensemble. The decoding procedure was performed for datasets K1-9 for the time intervals in table 4.2. Each combination of interval and method was decoded 100 times to obtain the distribution for each dataset. The boxes encompass the values between the 25th and 75 percentiles. The line and circle inside the box represent the median and mean values of the distribution, respectively. The fences limit the minimum and maximum values not considered outliers.

features of the total set, be effectively trained with the training set, and validated with the test set. In this analysis, it is also unquestionable that decoding with jPCA projections is significantly more efficient than using the entire neuronal ensemble

4.3.3 Decoding movement onset

This section aims to compare the rotational structure present in the neuronal activity during an interval after the object has been presented to the monkey and another interval right before the subject starts the actual movement. The goal of this analysis is to perceive if rotations are influenced by the timing of the action.

4.3.3.A Input to Support Vector Machine (SVM)

As previously shown, visuomotor information modulates neuronal activity during the whole duration of the task, after the subject starts planning said task, i.e. after object onset. However, dynamical structure might not be consistent during task planning and execution. Moreover, Krishna et al (2013) argues that the preparatory state (the neuronal state before movement onset) predicts the reaction time of the action, which motivates the following analysis. (9)

The hypothesis here is that, besides movement direction, the dynamics of neuronal activity may carry information relative to when the movement should start. To evaluate that, each trial is sampled during an interval early in the preparatory (visual) phase, and another interval immediately before the movement onset, as described in table 4.3. This way, each trial generates two observations.

Table 4.3: Intervals used in the decoding of movement onset. The column ≪Alignment≫ defines the event to which the starting and ending points are relative to.

Interval	Starting point (ms)	Ending point (ms)	Alignment
Not Onset	50	350	object onset
Onset	-300	0	lift hand

Similar to was done for decoding movement condition with the rotational dynamics, each feature corresponds to a time instant in each of the jPCA projections, as represented in fig. 4.18(a). As for the decoding with the entire neuronal ensemble, features correspond to the raw spike rates of all channels and all time points in the intervals, as represented in fig. 4.18(b). Unlike the previous analysis, there is no negative control to validate this experiment.

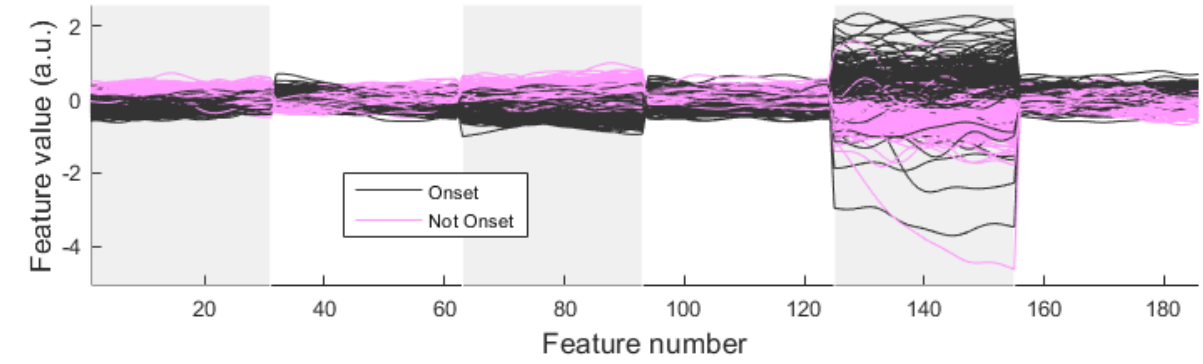

(a) Experiment

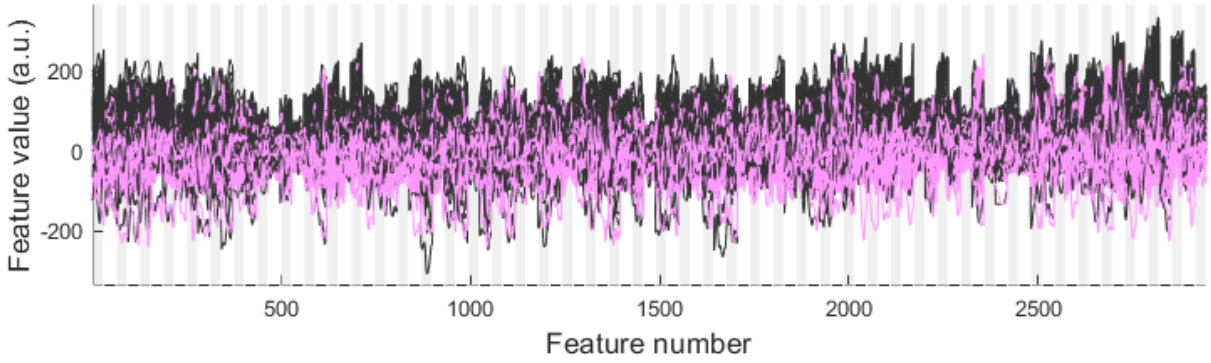

(b)

Figure 4.18: Example of the input observations used to decode and control the decoding of movement onset in monkey K2. **(b)** Features are the amplitude of the jPCA projections in each instant during the intervals in table 4.3. **(c)** Features are the rates of every recorded neuron after object onset at each instant in the same intervals as **(a)**. The grey-shaded areas identify the odd jPCA projections (for **(a)**) and channels (for **(b)**).

4.3.3.B Decoding accuracy

In order to obtain the distribution of the accuracy of decoding movement onset as represented in fig. 4.19, the SVM was trained 100 times for each dataset and for each feature design, using a ten-fold cross validation procedure to train and test the machine.

The input set of observations are designed to sample an 'onset' and a 'not onset' interval from each trial, and, thus, the chance level for this analysis is 50%. This value was highly surpassed with either feature design. In particular, decoding movement onset with jPCA projection had a distinctly superior performance that with the entire neuronal ensemble.

With the same motivation as before, a temporal analysis of the training duration was performed, and the results are plotted in fig. 4.20. For this application, decoding with jPCA projection also revealed to be more efficient than with the activity of each channel.

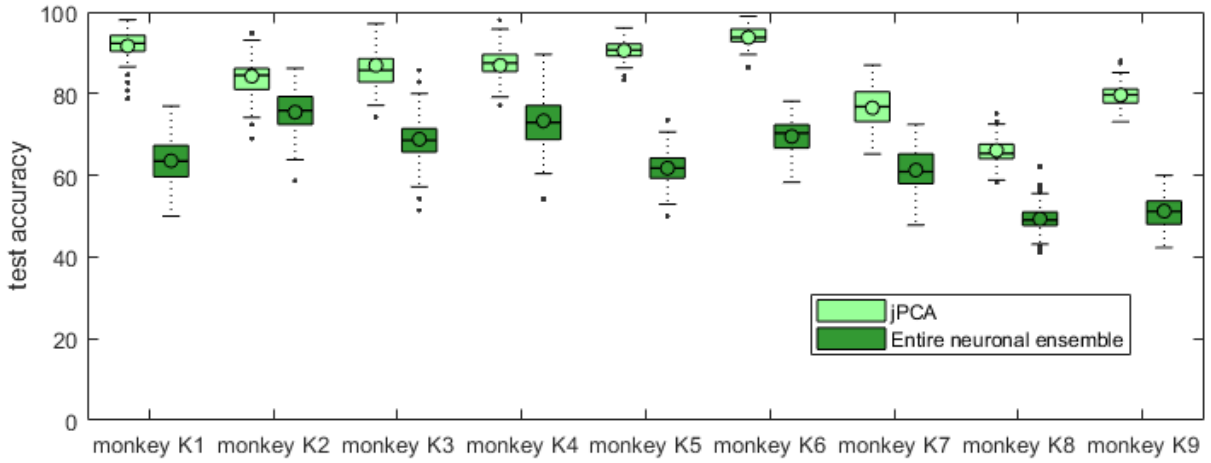

Figure 4.19: Accuracy of decoding movement direction with jPCA projections versus the entire neuronal ensemble. The decoding procedure was performed for datasets K1-9 for the time intervals in table 4.2. Each combination of interval and method was decoded 100 times to obtain the distribution for each dataset. The boxes encompass the values between the 25th and 75 percentiles. The line and circle inside the box represent the median and mean values of the distribution, respectively. The fences limit the minimum and maximum values not considered outliers.

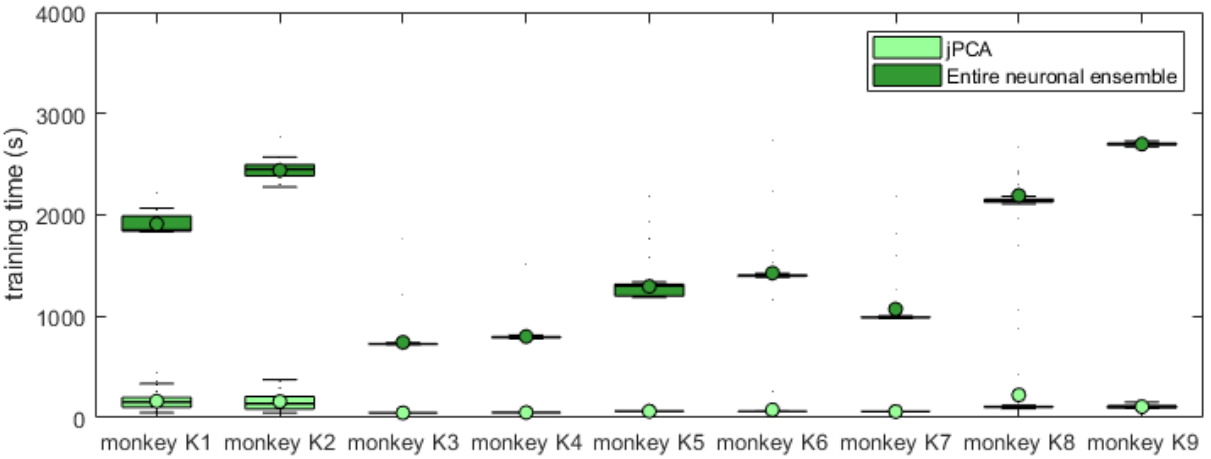

Figure 4.20: Time necessary to train the SVM for decoding movement direction with **(a)** jPCA projections versus **(b)** the entire neuronal ensemble. The decoding procedure was performed for datasets K1-9 labelled as in table 4.2. Each combination of interval and method was decoded 100 times to obtain the distribution for each dataset. The boxes encompass the values between the 25th and 75 percentiles. The line and circle inside the box represent the median and mean values of the distribution, respectively. The fences limit the minimum and maximum values not considered outliers.

Chapter 5

Discussion

In the interest of a good understanding and readability of the discussion, this chapter has a similar organisation to chapter 4. Each of the following sections aims to interpret the results within the framework of the present literature and to draw new conclusions based on the product of this research work.

5.1 Neuronal activity in PMd and PMv

Although across-channel average PSTHs often attenuate or misrepresent prominent features in the activity of individual channels, these plots are often an efficient way to distinguish between brain areas and their roles in the task at hands. Classically, PMd is associated with the planning of reaching tasks, whereas PMv has been implicated in the control of grasping movements. (7) And, in fact, by analysing the across-channel average PSTHs, namely that in fig. 4.3, the PMd population is active right after the target location of the object is known and the neurons in PMv spike at a higher rate during the movement itself, when the hand is shaping in preparation for the grasp.

When analysing the individual responses recorded in each of these 186 channels, 96 in each cortical region, a wide variety of responses came up. As expected for PMd, some channels recorded groups of neurons that responded more to visual stimuli, others that were more active during the reaching task itself, and even other channels that recorded neurons with both responses. PMv neurons in F5c are known to have motor- and visuomotor-dominant activity. (5) It, thus, comes as no surprise that most responsive neurons found described a stronger activity during the actual movement, with some of them showing discriminatory activity for different conditions. Furthermore, PMv has been associated with the preshaping of the hand for pre-

hension, which may explain the increase in firing rate at the start of the grasp movement.

Curiously, a few neurons in PMd, like those in fig. 4.4(d), showed distinct activity during the grasp of the target object, and some PMv neurons showed a visual-dominant responses which is not frequent for area F5c. (19) These observations suggest that this classical division of reach and grasp respectively in PMd and PMv is not strict, as already argued by Takahashi et al. (2017). (5) The presence of neuronal units in PMd and PMv that are active in distinct phases of motor planning and execution point towards a reach and grasp form a synergy. (5)

One would not expect a high discrimination between directions at an early stage in the planning phase. However, both areas showed high direction-sensitivity right after object onset. A likely explanation is that the monkey was well acquainted with the sequence of cues in the task and therefore exhibited early anticipatory activity.

5.2 Rotational structure in PMd and PMv dynamics

Overall, the MUA recorded in both cortical areas displayed rotational patterns in state space in all intervals. These rotational patterns were usually centred around the same point for all conditions, except for the third jPCA plane, where rotational patterns were more distinct between conditions. These results were consistent in all datasets analysed.

These rotational dynamics were not found underlying just a small portion of data variance in neither of the areas. In fact, the low-dimensional space described by the jPCA projections explained more than 74% and 73% of data variance for PMd and PMv, respectively, in all intervals and datasets. This slight difference may reflect a more varied neuronal response in PMv than in PMd. Nonetheless, these results surpass what Churchland et al. (2012) described when applying jPCA to model neuronal responses recorded in M1 and PMd ($50 - 70\%$). (8) This difference is possibly a consequence of them having a bigger number of unit isolations (our datasets only had 64 channels each), with more varied activity that requires more principal components to be accurately represented.

For the datasets and intervals analysed, the neuronal activity in both PMd and PMv seems to be equally well predicted by rotational dynamics. In particular, $R^2_{PMd} > 31\%$ and $R^2_{PMv} > 22\%$ during movement onset are of the same order of those found by Lara et al. (2018) in supplementary motor area (SMA) ($R^2_{SMA} > 0.22$), but much inferior to those found in M1 and PMv ($R^2_{M1\&PMd} > 0.76$) during a similar period. (56). Additionally, the rotational structure described explains the largest portion of the linear dynamics in both areas ($RGR_{PMd} > 73\%$ and $RGR_{PMv} > 78\%$, ignoring the outlier); a result superior to what has been described by Churchland et al. (2012) (8) (RGR slightly above 50% for a M1 and PMd population). These results

represent strong evidence that rotational structure underlie the neuronal trajectories in PMv populations as much as in PMd. It is, nonetheless, fundamental to analyse a larger set of neuronal recordings in order to reach a statistically significant conclusion.

There are two particularly interesting results worth noting. First, PMv does not show notorious response in the PSTHs during the visual phase; however, strong rotational activity is found in this same epoch. This comes to show how different the information represented in both plots is. Secondly, and most surprisingly, not only both areas displayed rotational dynamics even before the first cue - object onset -, but the goodness of fit of the model was superior in this interval. This, however, is linked to the fact that individual channels record very similar activity for any condition during this interval, as neuronal activity is still very homogeneous, which makes the neuronal signals more predictable, decreasing the fitting residuals - *see* eq. (3.6).

Unquestionably, the jPCA model fits well the activity in these two premotor areas, with rotational patterns that capture large portion of the data variance. However, these patterns could have appeared for trivial reasons or even by accident. The control analyses performed served the purpose of testing against that.

By disrupting the dynamical structure of data in three distinct ways while preserving the complexity of the response patterns, one can observe a quantitative loss in the rotations projected onto jPCA planes after the shuffling. This effect was best seen in fig. 4.10, where the original data is shown to have a distribution of measured angles peaked close to $+\pi/2$, whereas the shuffled data less circular structure. These results are concordant with previously reported data for macaque monkeys and humans. (2,8) Notably, the fraction of data that can be explained with rotational activity greatly decreases after shuffling the data, although the goodness of fit remains identical between PMd and PMv.

As expected, the shuffle controls did not remove the rotational structure from the data but reduced the consistency of any true rotations. Namely, the simulated data showed some clockwise rotations, when jPCA planes are defined to maximise anticlockwise rotation. For that reason, it is safe to conclude that the rotational structure in the original data is not simply a by-product of the technique, but rather, it is an intrinsic characteristic of the recorded signals. (2)

The ensuing control test determined that the quality of the dynamical fit of these datasets did not happen randomly. For most datasets and intervals, reorganising the dynamical trajectories of the across-trial populations in different conditions only improved the goodness of fit parameters, R^2 and RGR, less than 5% of the times for both premotor areas.

Another potential concern could be that different patches of cortex are more or less dynamical than others. Answering this hypothesis would require further experimentation, but this work can attest this is not a concern at a local scale, as the six datasets used have information

from different neuronal ensembles - *see* table 3.1 -, and yet all datasets display similar statistics.

These results highly suggest that the overall activity in both PMd and PMv is, indeed, very well predicted by rotational dynamics. These rotational dynamical patterns underlying the MUA are no surprise in PMd, as various studies have already reported and evaluated jPCA in this context. (2,8,62) The same cannot be said for PMv, whose rotational dynamics, as far as we know, have only been studied for LFPs or for single neurons, in both cases in a different context from this project. This work fills in the research gap in comparing the rotational dynamics that describe the MUA in both these areas. (63,64)

Rotational dynamics are a very common phenomenon in motor preparation, already studied in M1, PMd SMA and now PMv. (8,56) An interesting experiment to follow this work could be to determine where exactly neuronal activity gains these dynamical properties.

One hypothesis is that these rotational dynamics in PMv might be a reflection of inputs from PMd, as illustrated in fig. 2.2. The rationale behind this conjecture is that both areas work together during prehension tasks, planning and coordinating the reach (PMd) and the grasp (PMv), but their functional subdivision has been proved to not be as linear as previously thought. (5)

Another possible origin for these dynamical pattern is the parietal cortex itself, an area that processes visual information and send it to the premotor cortex. Should this hypothesis be confirmed, which hasn't been proved yet, then rotational dynamics would also be detected in other areas that are active when looking at an object but do not intervene in motor planning, such as the ITC. (15) If rotational dynamics are not found there, this might indicate that this structure is a signature of motor planning.

Mention that jPCA loses credibility for such a small nr of conditions

5.3 Unravelling the neuronal code in PMd

The rotations of the neuronal state are undoubtedly a robust feature of the physiological data recorded in the motor cortex. However, it is not immediately apparent what exactly these rotations reveal about movement in general, and reaching tasks in particular. (8)

Many research groups studying cortical control of arm movement are invested in orienting and timing the action. (9,55) Following the course of the recent research, this work uses the rotational structure to extract information about movement direction and onset. To our knowledge, no other study has used jPCA projections to directly cluster observation and decode neither movement direction nor timing.

Dynamical structure is typically studied on across-trial averaged data to cancel out recorded

noise and artefacts. But such an approach is far from ideal when classifying movement characteristics during a period of time, as the spike rate amplitude and timing of each channel in individual trials fluctuate, even if recorded under the same experimental condition. (64) The hypothesis underlying our approach is these minor variations in individual trials can be discarded and neuronal state of each trial can still be represented in a 6-dimensional state composed of only the major principal components, without prejudice to the motor information captured. This, however, naturally implies a decrease in the goodness of fit of the jPCA model.

The first conclusion that can be drawn from the decoding of movement direction is that, indeed, this motor parameter can be extracted from both jPCA projections and the entire neuron ensemble when that information is present, i.e. after object onset. On that note, as predicted, the precision of the classification for interval 1 - before object onset - is distributed around the chance level: 33%. For some datasets, this value might be only slightly above or slightly below that value, but that is entirely due to the constitution of the training and test sets, that are randomly generated from the entire set. This confirms the validity of the classification procedure, as observations are not correctly labelled more often than expected by chance.

In the intervals 2 and 3, both corresponding to epochs after object onset, the test accuracy was always superior to the chance level for both feature designs. These results should come as no surprise, as PMd is involved in preparing a motor program in response to a stimulus - object onset. (7,24)

A curious observation, and one of the most anticipated in this project, is that the decoding performance of movement direction using jPCA projection was undoubtedly superior to that using the entire neuron ensembles. These results are supported by several studies that report that low-dimensional dynamics can provide as accurate or even better reconstruction of 3D movement kinematics than approaches based on direct decoding from the entire neuron ensemble. (48,64) This observation suggests that different task conditions are characterised by a certain neuronal trajectory, and that this dynamical structure is a better descriptor of movement direction than the activity of the entire set of record neurons.

Furthermore, in both feature designs, there does not seem to exist any relation between the decoding performance in intervals 2 and 3. This results suggests that the flow-field of the neuronal activity in the low-dimensional (jPCA projections) and the high-dimensional (all channels) state-spaces does not necessarily gain or lose structure over time during the planning process. In other words, it looks like the subject is ready to execute the action early in the planning stage but is only waiting for the go cue. This result is particularly important, because it shows that non-cued behaviour, like real-life activities, can be very well predicted with dynamical rotations.

It is worth noting that the subject's arm was not constrained to a single pathway. In fact,

in some trials, before grasping the correct object, the monkey hesitates, and changes the direction of movement. This may be a explanation for some the differences in the neuronal trajectories.

On another note, decoding the onset of the action is highly dependent on differences in the trajectories during the planning of said motor command. Some experiments use the preparatory activity to predict the reaction time, while others rely on different data samples labelled onset or not onset to decode movement onset. (9,55) Implementing the first approach with an SVM, however, relies on the discretisation of the reaction time, which has the inherent problem that two neighbouring values might belong to different classes and other two more spaced values may be grouped together. Thus, the second method seems more fitting for the purpose of comparing the rotational structure before and during motor execution.

The accuracy of decoding movement onset with jPCA projection has a very narrow distribution with a distribution consistently centred much above the chance level, 50%. Not such good results were obtained when decoding with the entire neuronal ensemble, as some datasets had a similar performance to that of random classification.

Although few channels show increased activity right before movement onset, the major increase in spike rate occurs during the execution of a motor command. Thus, it would be expected that decoding with the neuronal ensemble would have a fairly good performance in this task; which was not verified. On the other hand, the rotational structure was not anticipated to be so distinct between the two interval as to allow such good characterisation of the activity in that period. These results suggest that there should be some condition-invariant dynamical structure that characterises the neuronal state at different epochs regardless of the experimental condition.

A possible explanation for this difference in rotational structure can be the go cue. It is important to rule out the hypothesis that the visual information associated with the go cue is not the cause of this dynamical organisation. This is especially relevant in real-life situations, as there is no cue involved and thus this form of decoding would not be effective.

One might be inclined to conclude that jPCA projections are best-suited for decoding movement onset than movement direction. But it is important to keep in mind that such comparison is not fair. In one case we are decoding a set with three labels, whereas in another case there are only two groups. Furthermore, it is not guaranteed that if we sampled each condition at a different time interval, the decoding performance would not be equally good of even better than that for decoding movement onset.

Both decoding procedures were successful and had visibly better results than when decoding with the entire neuronal set. Nonetheless, in a real-time application, the training and

decoding time are also crucial factors that may limit the application of a method, even if the machine is only training every other day. When training an SVM and choosing the best feature for that purpose, the less features there are to choose from, the faster the procedure will be. This reasoning explains why the SVM training time was so reduced when decoding with the jPCA projections compared to the entire neuronal set.

Chapter 6

Conclusion

In this study, we used an animal model to compare the rotational structure underlying the neuronal trajectories in PMd and PMv during a reach-to-grasp task and to understand what information might be encoded in the rotational dynamics of neuronal states in PMd. We demonstrated that both premotor regions are equally well-fit by a dynamical system, and that rotations comprehend almost the totality of the dynamical structure in both areas. Moreover, our results support the premise that movement parameters are more clearly evidenced in the temporal evolution of the neuronal population's state than in the activity of neuronal units. In particular, the neuronal trajectories in PMd during the planning phase were informative of the reach direction and predicted movement onset.

6.1 Functional properties of the premotor cortex

The difference in neuronal responses between the dorsal and ventral premotor cortices has strong implications for the development of novel BMIs based on the visuomotor activity of premotor neurons. In accordance with more recent studies, both premotor areas analysed showed a functional heterogeneity concerning the encoding of visual, reaching and grasping responses. Nonetheless, it should be noted that, in the datasets analysed, the neuronal preference was predominantly reaching in PMd and grasping in PMv, with PMd responding more strongly in the visual phase and PMv more active in the movement phase. Based on these results, one can reason that PMd is ideal not only for decoding reaching direction but for predicting the intended action before its execution. Neurons in PMv, on the other hand, seem to only distinguish direction after motion is initiated.

In contrast, the neuronal activity in both premotor areas seemed to be equally well cap-

tured by rotational dynamics in all epochs after object onset. Following the dynamical perspective of neuronal activity, this should be an argument supporting the role of both premotor areas in motor control during the whole duration of the reach-to-grasp task. This is an interesting finding, especially because it somehow contradicts what was expected from a representational point of view and raises two very important questions: where, in the visuomotor system, do neurons start showing rotational dynamics?, and is it be possible to decode the kinematics of the motor task from PMv during the planning phase?

As for the first question, two hypotheses arise: either dynamical rotational in PMv are a consequence of inputs from PMd, or this structure comes from an earlier stage of visual integration in the visuomotor system. To obtain a more definite answer, one could carry out further research in brain areas that are active when looking at an object but not involved in motor planning, like the ITC. The absence of rotational dynamics in this area could suggest that this neuronal population behaviour is a signature of motor planning.

Regarding the second point, it has been determined that PMv receives information about object shape, size and orientation from the PIP and object identity from theITC. Traditionally, this information has been thought to be used to adapt the shape of the hand for the prehension of the object. But there is no doubt that robust rotations are present in the neuronal trajectories of this premotor area even before movement onset, which leads us to believe that both premotor areas are able to discriminate between reach direction. Further research, however, would have to validate this belief and evaluate to what extent neuronal trajectories in PMv would be a significant addition to PMd, on what regards decoding reach direction. This exploration was the reason behind our intent to record simultaneous neuronal activity from PMd and PMv. But, as previously referred, this study was precluded due to technical contretemps.

Decoding with jPCA projections revealed to be consistently more accurate than using the entire neuronal set, as well as less time-consuming. This observation not only gives further support to a dynamical view of neuronal activity but also solidifies the relation between rotational dynamics and motor kinematics.

The results obtained reinforce the potential of controlling a motor BMI using premotor areas. Although only PMd data was decoded, the quality of the results suggests an equally good performance for PMv, at least in what regards the decoding of the grasping type. It has been described that both human and nonhuman primates can cognitively manipulate their activity in F5c and permit efficient decoding this way, which is a promising feature that can be an advantage for BMI control. (65,66) Moreover, continuous practice of a task has been reported to improve BMI's overall motor performance through cortical plastic adaptations, similar to mechanisms involved in the learning of new motor skills. (1,39,44,67,68)

6.2 Methodological considerations

An important issue that needs to be addressed is that we investigated the functional properties of the premotor cortex in a behaving rhesus monkey, where our ultimate goal is to develop novel BMI based on visuomotor activity for human use. Of all animal models used in brain research, the monkey brain is unquestionably the most similar to ours. However, evolution has provided the human brain with an elaborate neuronal network that has induced such reorganisation and expansion in premotor cortex, that studies have suggested that PMv in humans seems to be essential for speech processing. (69) These findings may complicate the implantation of the array in the human premotor cortex, since we are mostly interested in populations that are active during reach-to-grasp tasks.

Although jPCA is a simple method with comproved results, it does have some drawbacks that are addressed by another method, HDR. (56, 70) While jPCA removes the time-varying cross-condition mean in a very unprincipled way, HDR isolates the condition-invariant and condition-specific structure by projection onto orthogonal dimensions. (56) This method should be applied if one is concerned that weak rotational structure might be exaggerated. Other than that, both methods have been described to produce very similar results for motor cortex data. (70)

An additional concern in this project is related to the feature design when decoding motor parameters. SVMs are not designed to account for the temporal structure of variables, so an alternative is to treat each time point of each variable as a different feature. This approach, however, may compromise the decoding performance, begging the question of whether the decoding performance would be different were the temporal structure considered in the decoding process. In our point of view, for the purpose of offline decoding, this approach, although unrealistic, shoulf have a minor effect on the decoding accuracy, because each timestamp is compared against the same instant in other trials. Nonetheless, there are dynamical artificial network architectures, like the Long Short-Term Memory (LSTM), that can decode with sequences of data. Although not included in this thesis, an attempt with this algorithm was made, but the results were not improved relative to the present approach.

Another important decision in BMI control with intracortical signals is exactly which signal to use. In this project, MUA was used because of its stability over time. However, SUA are thought to contain the most selective information, which would be very appealing, were it not for its instability and the elaborate methods envolved in spike sorting. As a third alternative, it could be interesting to evaluate the performance of LFPs dynamics in decoding various motor parameters.

6.3 Future perspectives for real-time BMI control

The interesting findings we obtained should encourage us to continue researching in premotor areas with the purpose of creating and improving BMIs that can be controlled with visuomotor information. This technology could benefit patients with some level of motor-impairment (i.e. paralysis, locked-in syndrome, amyotrophic lateral sclerosis). A future goal of research in this field is to enhance BMI performance with more robust and energy-efficient decoding algorithms and the implantation of wireless arrays with hundreds of electrodes, to increase accuracy.

One of the first improvements that can be done regards the degrees of freedom that can be controlled with the visuomotor activity. We only looked at three different reach direction, but this number could be increased. Moreover, the objects could be placed at different distances from the subject, and different sizes and shapes of objects could be used - i.e. spheres, cubes, handles. This would allow the study of a combination of reach kinematics with different types of grips - i.e. power grip, precision grip, key grip -, allowing a more detailed comprehension of the structure of neuronal dynamics in both premotor areas in various conditions.

Should one want to perform online decoding of movement kinematics in a continuous space, this simplistic approach would not have been possible. In that case, further parameterisation would be necessary to represent the continuous space as a function of the jPCA projections. (64) Recent studies suggest the use of Kalman filters for decoding hand kinematics in a continuous space and stress the potential of this decoder for real-time applications. (64,71)

To extend this fundamental knowledge to the clinical setting of a BMI, experiments could be carried out in which a monkey has to move an avatar arm using its own neuronal activity. This way, one could not only study the plastic adaptations that occur in the premotor cortex when the subject changes their body representation in the brain, but also understand what happens at the neuronal level in this area when the monkey makes a mistake and corrects for it.

The neuronal activity varies with the position of the target relative to the direction of gaze. (4) For this reason, experimental setups often require the monkey to fix their eyes on a light, without moving, in order to assure a standardisation of the experimental conditions. In real life, however, our eyes are constantly performing visual exploration. Future studies should also account for the implications that saccades have on the dynamical structure of neuronal trajectories and on the decoding of kinematic parameters.

Currently, the Utah array is the only microelectrode implant approved for human use by the FDA. However, the insertion of these arrays is usually damaging, causing the formation of scar tissue around the array, which is associated with signal attenuation. (1) Further research is, thus, necessary to create novel implants and insertion methods that reduce fibrous encapsulation.

Bibliography

(1) M. A. Lebedev and M. A. Nicolelis, "Brain-machine interfaces: From basic science to neuroprostheses and neurorehabilitation," *Physiological reviews*, vol. 97, no. 2, pp. 767–837, 2017.

(2) C. Pandarinath, V. Gilja, C. H. Blabe, P. Nuyujukian, A. A. Sarma, B. L. Sorice, E. N. Eskandar, L. R. Hochberg, J. M. Henderson, and K. V. Shenoy, "Neural population dynamics in human motor cortex during movements in people with als," *Elife*, vol. 4, p. e07436, 2015.

(3) M. A. Dimyan and L. G. Cohen, "Neuroplasticity in the context of motor rehabilitation after stroke," *Nature Reviews Neurology*, vol. 7, no. 2, p. 76, 2011.

(4) E. R. Kandel, J. H. Schwartz, T. M. Jessel, S. A. Siegelbaum, and A. J. Hudspeth, *Principles of Neural Science, Fifth Edition*. The McGraw-Hill Companies, Inc, 2013, vol. 2.

(5) K. Takahashi, M. D. Best, N. Huh, K. A. Brown, A. A. Tobaa, and N. G. Hatsopoulos, "Encoding of both reaching and grasping kinematics in dorsal and ventral premotor cortices," *Journal of Neuroscience*, vol. 37, no. 7, pp. 1733–1746, 2017.

(6) P. A. Chouinard, "Different roles of pmv and pmd during object lifting," *Journal of Neuroscience*, vol. 26, no. 24, pp. 6397–6398, 2006.

(7) E. Hoshi and J. Tanji, "Differential involvement of neurons in the dorsal and ventral premotor cortex during processing of visual signals for action planning," *Journal of neurophysiology*, vol. 95, no. 6, pp. 3596–3616, 2006.

(8) M. M. Churchland, J. P. Cunningham, M. T. Kaufman, J. D. Foster, P. Nuyujukian, S. I. Ryu, and K. V. Shenoy, "Neural population dynamics during reaching," *Nature*, vol. 487, no. 7405, p. 51, 2012.

(9) K. V. Shenoy, M. Sahani, and M. M. Churchland, "Cortical control of arm movements: a dynamical systems perspective," *Annual review of neuroscience*, vol. 36, pp. 337–359, 2013.

(10) J. Krakauer and C. Ghez, "Voluntary movement," *Principles of neural science*, vol. 4, pp. 756–781, 2000.

(11) A. W. Campbell, *Histological studies on the localisation of cerebral function*. University Press, 1905.

(12) P. Janssen, B.-E. Verhoef, and E. Premereur, "Functional interactions between the macaque dorsal and ventral visual pathways during three-dimensional object vision," *Cortex*, vol. 98, pp. 218–227, 2018.

(13) F. Katsuki and C. Constantinidis, "Bottom-up and top-down attention," *The Neuroscientist : a review journal bringing neurobiology, neurology and psychiatry*, vol. 20, 12 2013.

(14) S. staff. (2009) Visuomotor skills. (Online). Available: https://www.aboutkidshealth.ca/Article?contentid=1879&language=English

(15) A. H. Fagg and M. A. Arbib, "Modeling parietal–premotor interactions in primate control of grasping," *Neural Networks*, vol. 11, no. 7-8, pp. 1277–1303, 1998.

(16) E. Oztop, "Modeling the mirror: grasp learning and action recognition," Ph.D. dissertation, University of Southern California, 2002.

(17) G. Rizzolatti and M. Matelli, "Two different streams form the dorsal visual system: anatomy and functions," *Experimental brain research*, vol. 153, no. 2, pp. 146–157, 2003.

(18) G. Rizzolatti, R. Camarda, L. Fogassi, M. Gentilucci, G. Luppino, and M. Matelli, "Functional organization of inferior area 6 in the macaque monkey," *Experimental brain research*, vol. 71, no. 3, pp. 491–507, 1988.

(19) A. Murata, W. Wen, and H. Asama, "The body and objects represented in the ventral stream of the parieto-premotor network," *Neuroscience research*, vol. 104, pp. 4–15, 2016.

(20) M. Umiltà, I. Intskirveli, F. Grammont, M. Rochat, F. Caruana, A. Jezzini, V. Gallese, G. Rizzolatti *et al.*, "When pliers become fingers in the monkey motor system," *Proceedings of the National Academy of Sciences*, vol. 105, no. 6, pp. 2209–2213, 2008.

(21) G. Rizzolatti and C. Sinigaglia, "The functional role of the parieto-frontal mirror circuit: interpretations and misinterpretations," *Nature reviews neuroscience*, vol. 11, no. 4, p. 264, 2010.

(22) T. Theys, M. C. Romero, J. van Loon, and P. Janssen, "Shape representations in the primate dorsal visual stream," *Frontiers in computational neuroscience*, vol. 9, p. 43, 2015.

(23) G. Rizzolatti, L. Cattaneo, M. Fabbri-Destro, and S. Rozzi, "Cortical mechanisms underlying the organization of goal-directed actions and mirror neuron-based action understanding," *Physiological reviews*, vol. 94, no. 2, pp. 655–706, 2014.

(24) M. J. Grol, J. Majdandzi, K. E. Stephan, L. Verhagen, H. C. Dijkerman, H. Bekkering, F. A. Verstraten, and I. Toni, "Parieto-frontal connectivity during visually guided grasping," *Journal of Neuroscience*, vol. 27, no. 44, pp. 11 877–11 887, 2007.

(25) P. Cisek and J. F. Kalaska, "Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action," *Neuron*, vol. 45, no. 5, pp. 801–814, 2005.

(26) M. M. Churchland, M. Y. Byron, S. I. Ryu, G. Santhanam, and K. V. Shenoy, "Neural variability in premotor cortex provides a signature of motor preparation," *Journal of Neuroscience*, vol. 26, no. 14, pp. 3697–3712, 2006.

(27) K. Kurata, "Premotor cortex of monkeys: set-and movement-related activity reflecting amplitude and direction of wrist movements," *Journal of Neurophysiology*, vol. 69, no. 1, pp. 187–200, 1993.

(28) E. López-Larraz, A. Sarasola-Sanz, N. Irastorza-Landa, N. Birbaumer, and A. Ramos-Murguialday, "Brain-machine interfaces for rehabilitation in stroke: A review," *NeuroRehabilitation*, vol. 43, no. 1, pp. 77–97, 2018.

(29) E. G. Jones, "Golgi, cajal and the neuron doctrine," *Journal of the History of the Neurosciences*, vol. 8, no. 2, pp. 170–178, 1999.

(30) D. Hebb, "The organization of behavior. emphnew york," 1949.

(31) P. G. Patil, J. M. Carmena, M. A. Nicolelis, and D. A. Turner, "Ensemble recordings of human subcortical neurons as a source of motor control signals for a brain-machine interface," *Neurosurgery*, vol. 55, no. 1, pp. 27–38, 2004.

(32) N. Fitzsimmons, M. Lebedev, I. Peikon, and M. A. Nicolelis, "Extracting kinematic parameters for monkey bipedal walking from cortical neuronal ensemble activity," *Frontiers in integrative neuroscience*, vol. 3, p. 3, 2009.

(33) J. E. O'Doherty, M. A. Lebedev, P. J. Ifft, K. Z. Zhuang, S. Shokur, H. Bleuler, and M. A. Nicolelis, "Active tactile exploration using a brain–machine–brain interface," *Nature*, vol. 479, no. 7372, p. 228, 2011.

(34) M. W. Slutzky and R. D. Flint, "Physiological properties of brain-machine interface input signals," *Journal of neurophysiology*, vol. 118, no. 2, pp. 1329–1343, 2017.

(35) J. An, S. H. Jin, S. H. Lee, G. Jang, B. Abibullaev, H. Lee, and J.-I. Moon, "Cortical activation pattern for grasping during observation, imagery, execution, fes, and observation-fes integrated bci: An fnirs pilot study," in *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2013, pp. 6345–6348.

(36) K. B. Clancy, A. C. Koralek, R. M. Costa, D. E. Feldman, and J. M. Carmena, "Volitional modulation of optically recorded calcium signals during neuroprosthetic learning," *Nature neuroscience*, vol. 17, no. 6, p. 807, 2014.

(37) E. Buch, C. Weber, L. G. Cohen, C. Braun, M. A. Dimyan, T. Ard, J. Mellinger, A. Caria, S. Soekadar, A. Fourkas *et al.*, "Think to move: a neuromagnetic brain-computer interface (bci) system for chronic stroke," *Stroke*, vol. 39, no. 3, pp. 910–917, 2008.

(38) D. T. Bundy, E. Zellmer, C. M. Gaona, M. Sharma, N. Szrama, C. Hacker, Z. V. Freudenburg, A. Daitch, D. W. Moran, and E. C. Leuthardt, "Characterization of the effects of the human dura on macro-and micro-electrocorticographic recordings," *Journal of neural engineering*, vol. 11, no. 1, p. 016006, 2014.

(39) J. M. Carmena, M. A. Lebedev, R. E. Crist, J. E. O'Doherty, D. M. Santucci, D. F. Dimitrov, P. G. Patil, C. S. Henriquez, and M. A. Nicolelis, "Learning to control a brain–machine interface for reaching and grasping by primates," *PLoS biology*, vol. 1, no. 2, p. e42, 2003.

(40) J. Krüger, F. Caruana, G. Rizzolatti *et al.*, "Seven years of recording from monkey cortex with a chronically implanted multiple microelectrode," *Frontiers in neuroengineering*, vol. 3, p. 6, 2010.

(41) S.-J. Kim, S. C Manyam, D. J Warren, and R. A Normann, "Electrophysiological mapping of cat primary auditory cortex with multielectrode arrays," *Annals of biomedical engineering*, vol. 34, pp. 300–9, 03 2006.

(42) C. A. Chestek, V. Gilja, P. Nuyujukian, J. D. Foster, J. M. Fan, M. T. Kaufman, M. M. Churchland, Z. Rivera-Alvidrez, J. P. Cunningham, S. I. Ryu *et al.*, "Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex," *Journal of neural engineering*, vol. 8, no. 4, p. 045005, 2011.

(43) E. V. Evarts, "Pyramidal tract activity associated with a conditioned hand movement in the monkey." *Journal of Neurophysiology*, vol. 29, no. 6, pp. 1011–1027, 1966.

(44) M. A. Nicolelis and M. A. Lebedev, "Principles of neural ensemble physiology underlying the operation of brain–machine interfaces," *Nature reviews neuroscience*, vol. 10, no. 7, p. 530, 2009.

(45) M. D. Best, "The involvement of premotor cortex in executing reach to grasp movements," University of Chicago, Tech. Rep., 2016.

(46) A. Leonardo, "Degenerate coding in neural systems," *Journal of Comparative Physiology A*, vol. 191, no. 11, pp. 995–1010, 2005.

(47) S. Maren, K. L. Phan, and I. Liberzon, "The contextual brain: implications for fear conditioning, extinction and psychopathology," *Nature reviews neuroscience*, vol. 14, no. 6, p. 417, 2013.

(48) J. A. Michaels, B. Dann, and H. Scherberger, "Neural population dynamics during reaching are better explained by a dynamical system than representational tuning," *PLoS computational biology*, vol. 12, no. 11, p. e1005175, 2016.

(49) D. Sussillo, M. M. Churchland, M. T. Kaufman, and K. V. Shenoy, "A neural network that finds a naturalistic solution for the production of muscle activity," *Nature neuroscience*, vol. 18, no. 7, p. 1025, 2015.

(50) T. Pistohl, A. Schulze-Bonhage, A. Aertsen, C. Mehring, and T. Ball, "Decoding natural grasp types from human ecog," *Neuroimage*, vol. 59, no. 1, pp. 248–260, 2012.

(51) D. M. Wolpert and Z. Ghahramani, "Computational principles of movement neuroscience," *Nature neuroscience*, vol. 3, no. 11s, p. 1212, 2000.

(52) D. Garrett, D. A. Peterson, C. W. Anderson, and M. H. Thaut, "Comparison of linear, nonlinear, and feature selection methods for eeg signal classification," *IEEE Transactions on neural systems and rehabilitation engineering*, vol. 11, no. 2, pp. 141–144, 2003.

(53) E. Hortal, D. Planelles, A. Costa, E. Iáñez, A. Úbeda, J. M. Azorín, and E. Fernández, "Svm-based brain–machine interface for controlling a robot arm through four mental tasks," *Neurocomputing*, vol. 151, pp. 116–121, 2015.

(54) T. Sakurada, T. Kawase, K. Takano, T. Komatsu, and K. Kansaku, "A bmi-based occupational therapy assist suit: asynchronous control by ssvep," *Frontiers in neuroscience*, vol. 7, p. 172, 2013.

(55) Z. Wang, A. Gunduz, P. Brunner, A. L. Ritaccio, Q. Ji, and G. Schalk, "Decoding onset and direction of movements using electrocorticographic (ecog) signals in humans," *Frontiers in neuroengineering*, vol. 5, p. 15, 2012.

(56) A. H. Lara, J. P. Cunningham, and M. M. Churchland, "Different population dynamics in the supplementary motor area and motor cortex during reaching," *Nature communications*, vol. 9, no. 1, p. 2754, 2018.

(57) MathWorks. (2008) sequentialfs. (Online). Available: https://nl.mathworks.com/help/stats/sequentialfs.html

(58) G. Zheng, Z. Qian, Q. Yang, C. Wei, L. Xie, Y. Zhu, and Y. Li, "The combination approach of svm and ecoc for powerful identification and classification of transcription factor," *BMC bioinformatics*, vol. 9, no. 1, p. 282, 2008.

(59) MathWorks. (2014) fitcecoc. (Online). Available: https://nl.mathworks.com/help/stats/fitcecoc.html

(60) ——. (2014) fitcsvm. (Online). Available: https://nl.mathworks.com/help/stats/fitcsvm.html

(61) R. Brette, "Philosophy of the spike: rate-based vs. spike-based theories of the brain," *Frontiers in systems neuroscience*, vol. 9, p. 151, 2015.

(62) C. Pandarinath, D. J. O'Shea, J. Collins, R. Jozefowicz, S. D. Stavisky, J. C. Kao, E. M. Trautmann, M. T. Kaufman, S. I. Ryu, L. R. Hochberg *et al.*, "Inferring single-trial neural population dynamics using sequential auto-encoders," *Nature methods*, p. 1, 2018.

(63) T. M. Hall, F. de Carvalho, and A. Jackson, "A common structure underlies low-frequency cortical dynamics in movement, sleep, and sedation," *Neuron*, vol. 83, no. 5, pp. 1185–1199, 2014.

(64) M. Aghagolzadeh and W. Truccolo, "Inference and decoding of motor cortex low-dimensional dynamics via latent state-space models," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 24, no. 2, pp. 272–282, 2015.

(65) J. Simeral, S.-P. Kim, M. Black, J. Donoghue, and L. Hochberg, "Neural control of cursor trajectory and click by a human with tetraplegia 1000 days after implant of an intracortical microelectrode array," *Journal of neural engineering*, vol. 8, no. 2, p. 025027, 2011.

(66) M. D. Serruya, N. G. Hatsopoulos, L. Paninski, M. R. Fellows, and J. P. Donoghue, "Brain-machine interface: Instant neural control of a movement signal," *Nature*, vol. 416, no. 6877, p. 141, 2002.

(67) K. Ganguly and J. M. Carmena, "Emergence of a stable cortical map for neuroprosthetic control," *PLoS biology*, vol. 7, no. 7, p. e1000153, 2009.

(68) P. J. Ifft, S. Shokur, Z. Li, M. A. Lebedev, and M. A. Nicolelis, "A brain-machine interface enables bimanual arm movements in monkeys," *Science translational medicine*, vol. 5, no. 210, pp. 210ra154–210ra154, 2013.

(69) I. G. Meister, S. M. Wilson, C. Deblieck, A. D. Wu, and M. Iacoboni, "The essential role of premotor cortex in speech perception," *Current Biology*, vol. 17, no. 19, pp. 1692–1696, 2007.

(70) Churchland lab: code. (Online). Available: https://churchland.zuckermaninstitute.columbia.edu/content/code

(71) W. Wu, Y. Gao, E. Bienenstock, J. P. Donoghue, and M. J. Black, "Bayesian population decoding of motor cortical activity using a kalman filter," *Neural computation*, vol. 18, no. 1, pp. 80–118, 2006.

(72) M. Ringnér, "What is principal component analysis?" *Nature biotechnology*, vol. 26, no. 3, p. 303, 2008.

Appendix A

Geometrical interpretation of jPCA

This section aims to provide a more geometrical understanding of the algorithms used to reduce state-space dimensionality, PCA, and to find planes with significant rotational structure, jPCA. An extensive algebraic description of PCA is deemed unnecessary, as this is an established method, and bears no influence on the functioning of the core method of this work - jPCA -, already carefully described in section 3.2. For a better understanding, the following explanation seeks to keep the notation coherent with the aforementioned section.

Considering first a set of high-dimensional data, X, represented in fig. A.1, the first step in the algorithm is to reduce its dimensionality. For that application, standard PCA is performed.

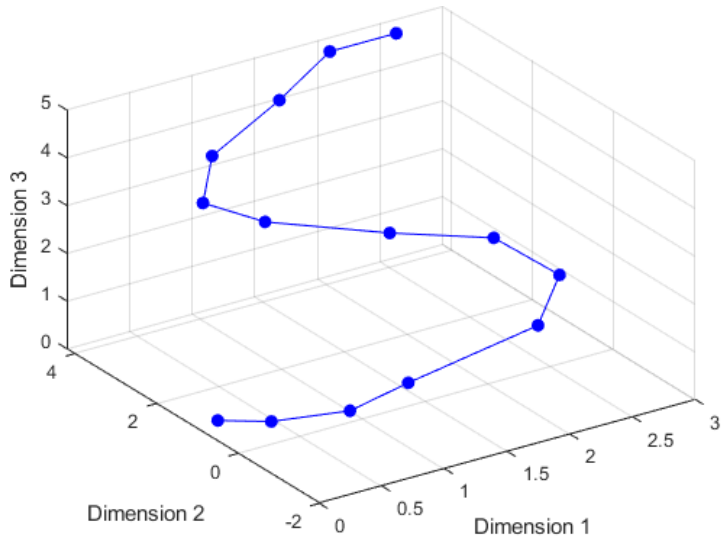

Figure A.1: 3-Dimensional artificial data, X, used in the PCA and jPCA explanations

PCA uses an orthogonal transformation, such that the first principal component retains the most variance possible for a single vector, and every following component captures the most

variance possible, ensuring orthogonality to all preceding components. (72)

In the 3-dimensional state-space illustrated in fig. A.2(a), both the artificial data (in blue) and the low-dimensional data (in orange) are represented. For this representation of the low-dimensional data, only the first two principal components (represented in black and red), PCs, were used. The top plot shows a clear resemblance in the trajectories of both sets of data. The bottom one clearly reveals that the low-dimensional data are spread on the surface defined by PC1 and PC2 (aligned with the orange line). It is worth noting that PCA centres the newfound axis on the mean of the original data. Figure A.2(b) shows the projection of the original data onto the plane defined by the first two Principal Components (PCs) on a 2-dimensional surface.

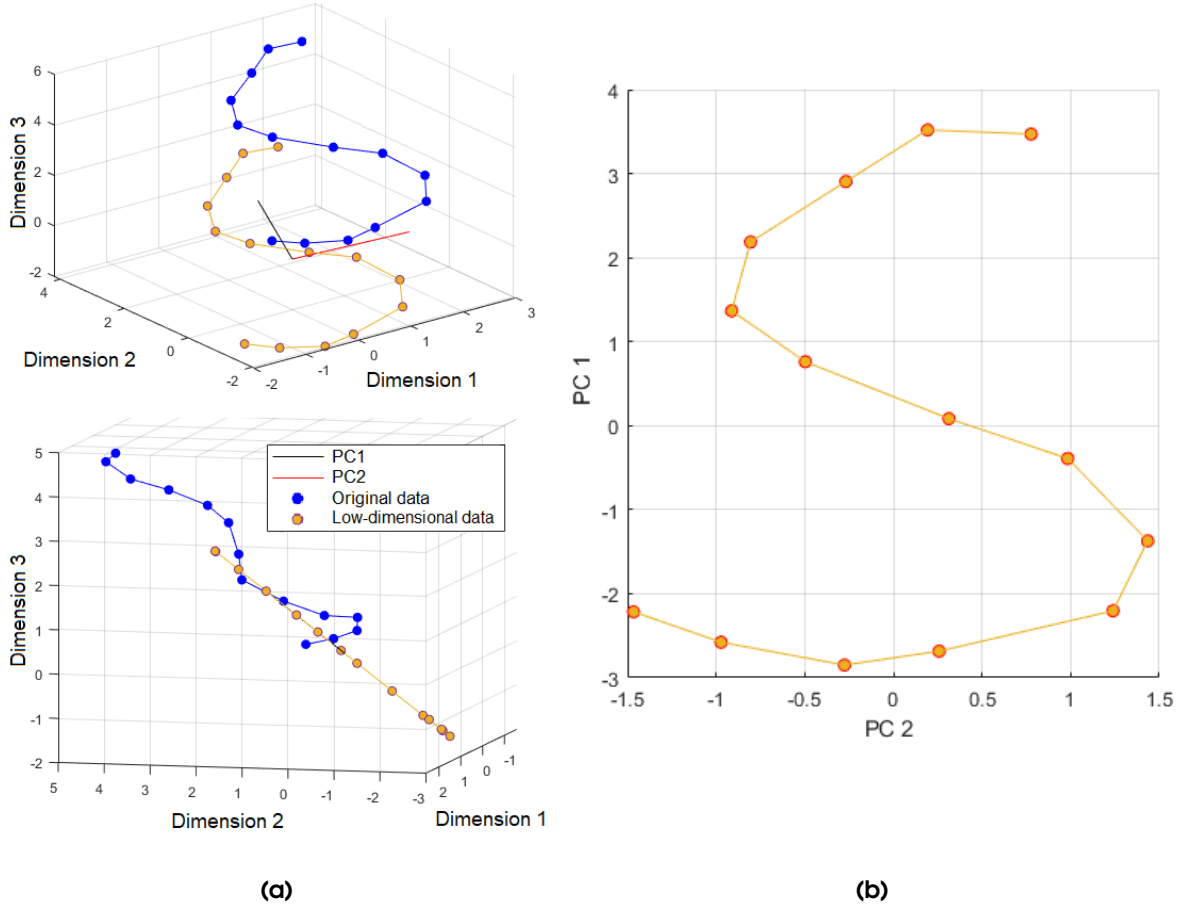

Figure A.2: Result of performing PCA on the artificial data. **(a)** Two views of the 3D artificial data projected (in blue) onto the first two principal components (in orange). The first and second principal components are represented in black and red, respectively. **(b)** projection of the 3D data onto the first PCA plane, defined by the first two principal components.

In normal circumstances, jPCA would use the low-dimensional data to find a plane with significant rotational structure. However, the data after PCA is already a planar. Alternatively, the same dataset, X, will be used to aid the geometrical understanding of jPCA.

Section 3.2 introduced the decomposition of the general linear dynamical into a symmetric and a skew-symmetrical (referred to as only skew) parts, describing scaling and rotations of data, respectively.

Any matrix, M, describing a linear dynamical system can be represented as $M = VDV^{-1}$, where the i-th column of the square-matrix V is the eigenvector q_i of M, and D is a diagonal matrix whose i-th element is the eigenvalue corresponding to q_i. Thus, any linear dynamical systems, regardless of the constraints to M, can be described as in eq. (A.1).

$$\begin{aligned} \dot{X} &= XM \\ \dot{X} &= XVDV^{-1} \\ \dot{X}V &= (XV)D \\ \dot{U} &= UD \end{aligned} \tag{A.1}$$

Through diagonalisation of the matrix that describes the dynamical system, it becomes much more obvious that each component of the state variation (each entry in a row), $\dot{U} = \dot{X}V$, is a mere scaling of the said state, $U = XV$, if both are in the eigenvector basis. The scaling factor, i.e., the eigenvalue, can be a real or purely imaginary number, in the case of symmetric and skew matrices M, respectively, and that is where the main difference between both systems resides.

As the eigenvectors of M_{skew} are complex, the representation of the data components in this space, U will also be complex. Given an initial component of the data in this basis, $u_1(t_0)$, scaling it by a purely imaginary number will result in a rotation in the complex plane corresponding to $\dot{u_1}(t = 0)$. (It may be helpful to remember that every complex vector multiplied by i is rotated by $\pi/2\,\mathrm{rad}$ in the complex plane.) As a consequence, the same component, u_1, will be rotated in next time instant, as $u_1(t_1) = u_1(t_0) + \dot{u_1}(t_1)$. For the real-valued eigenvalues of M_{sym}, each component and its derivative are always along the same direction.

To illustrate the previous explanation, M_{skew} was computed following the jPCA algorithm described in section 3.2, and M_{sym} was obtained by defintion, $M_{sym} = (M + M^T)$. Figure A.3 shows the artificial data projected onto the two eigenvectors with the biggest eigenvalues for each matrix. Because the two eigenvectors of M_{skew} with the most eigenvalue magnitude are complex conjugate, the representation of the data in the eigenvector basis (fig. A.3(a)) is symmetric to the real axis, and follows a non-linear trajectory in the complex space. Contrarily, the representation of the data components in the basis of the first two eigenvectors of M_{sym} (fig. A.3(b)) is constantly real-valued.

The next step is to find a surface that represents expansions/contractions of data, for the

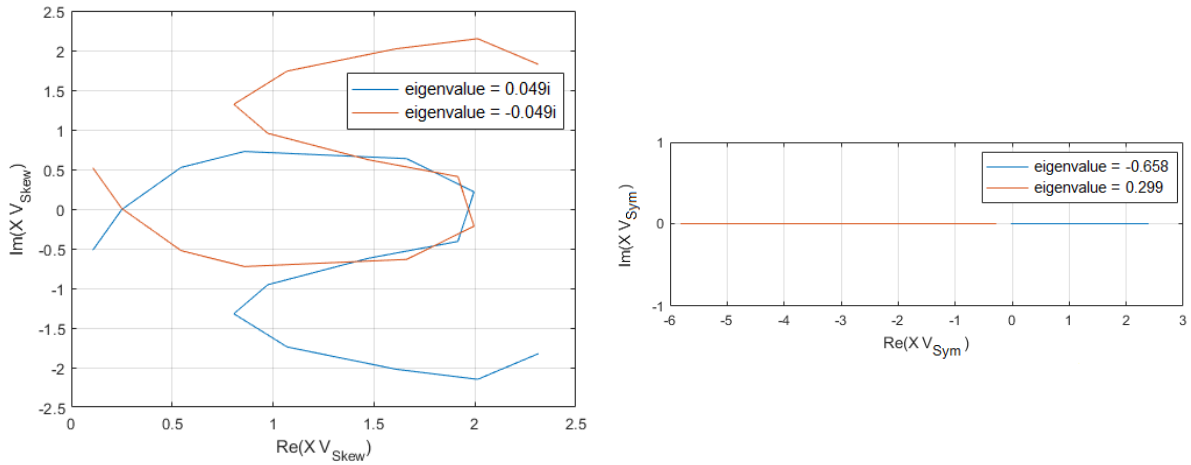

(a) **(b)**

Figure A.3: Projection of the artificial data, X, onto the eigenvectors with the biggest eigenvalue norm: **(a)** V_{skew} are eigenvectors of M_{skew}; and **(b)** V_{sym} are eigenvectors of M_{sym}.

case of M_{sym}, and significant rotational structure, for the case of M_{skew}.

For the first case, plotting the two components of the data projected onto the two (orthogonal) eigenvectors can be easily achieved. In fact, fig. A.4(a) displays the plane defined by the two eigenvectors, named symPC1 and symPC2 by analogy to the jPCA projections (jPCs). Figure A.4(b) shows a view of the original data that has approximately the same structure as the projection on the plane found with M_{sym}.

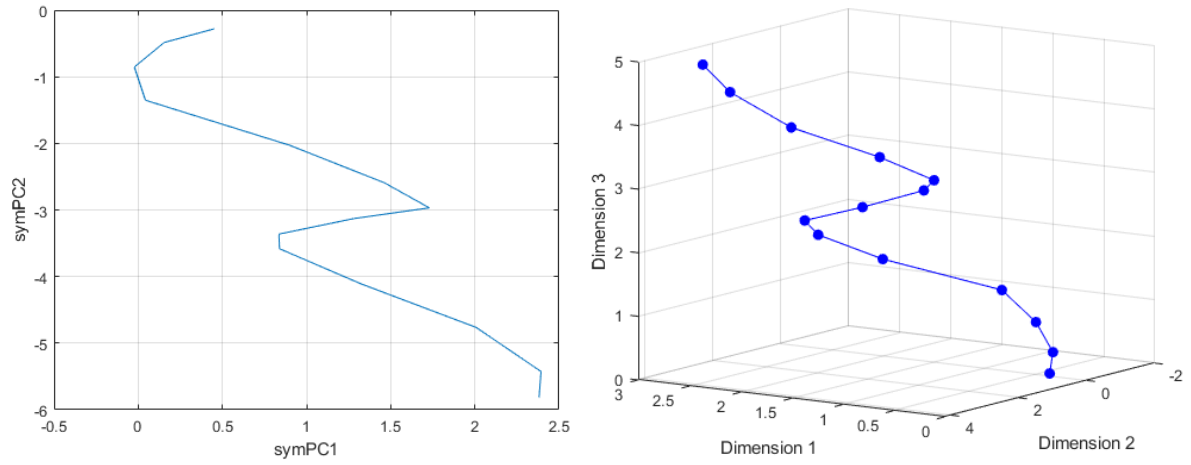

(a) **(b)**

Figure A.4: Projection of data onto a plane defined by the eigenvectors of a symmetric summary matrix, M_{sym}: **(a)** projection of the artificial data onto a plane defined by the two eigenvectors in V_{sym} with the biggest eigenvalue norm; **(b)** view of the artificial data, X, that approximately corresponds to the plane represented in **(a)**.

For the case of M_{skew}, the two projections cannot be easily plotted against each other. Despite their orthogonality, each component is 2-dimensional (has both real and imaginary parts), which would generate a 4-dimensional plot. Alternatively, jPCA creates a 2-vector base corresponding to the imaginary and real parts of the eigenvector with the positive eigenvalue. This basis is, then, rotated so that the data is mostly spread along the horizontal axis and the trajectory displayed rotates anticlockwise. Figure A.5(a) shows the projection of the data onto this plane. In this example, despite a there being clockwise rotation, the majority of the points have a derivative leading to an anticlockwise rotation. Figure A.5(b) displays a view of the original data that matches the plane found by jPCA as being the one with the most significant rotations. This comes to show that jPCA does not alter the data, but simply rotates it to make its rotational architecture more obvious.

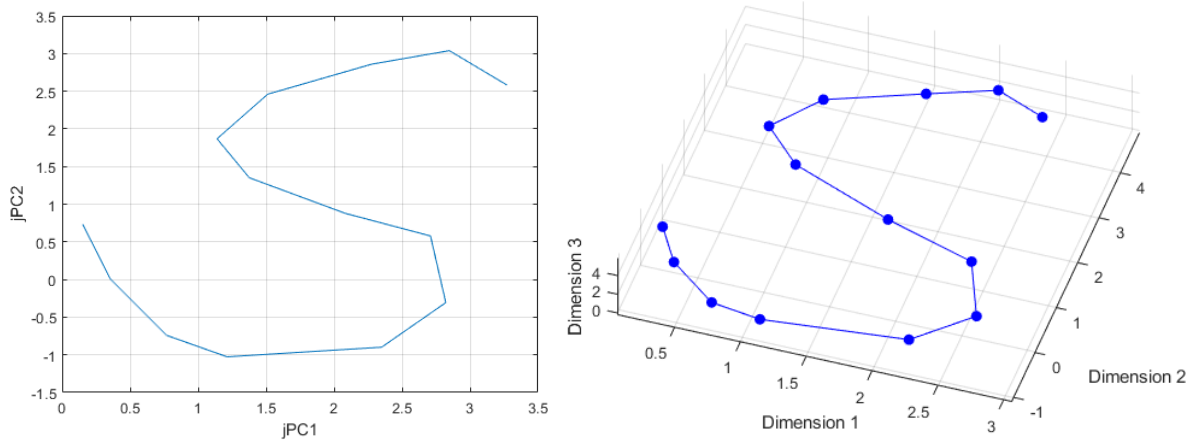

(a) **(b)**

Figure A.5: Projection of data onto a jPCA plane defined by the eigenvectors of a skew-symmetric summary matrix, M_{skew}: **(a)** projection of the artificial data onto a plane defined by the two jPCs, where each jPC is the combination of the two complex conjugate eigenvectors in V_{sym}; **(b)** View of the artificial data, X, that approximately corresponds to the jPCA plane in **(a)** and found by jPCA.

Appendix B

State-space representation of disrupted data

This chapter follows from section 4.2.3.A, where a set of simulated data was used to assess the power of jPCA.

Figure B.1 and fig. B.2 illustrate the PSTHs for the original and shuffled data for two channels in PMd and PMv, respectively. It is visible that the plots start differing after the time-point at 100 ms; afterwards, the behaviour of each condition evolves differently according to the shuffle control.

The jPCA projections of the original and simulated PMd activity are represented in fig. B.3 (Cont.). The original data, in fig. B.3(a), displays strong rotations in 89% of the data variance captured, especially in the first two jPCA planes. The shuffle control 1 causes most of the robustness of rotational structure to be lost. In fig. B.3(b), only the first jPCA plane presents curved trajectories, which corresponds to 13% of data variance. As depicted in fig. B.3(c), the individual trajectories in shuffle control 2 present almost no rotational structure, and the existing rotations are not robust; also, the direction of the curvature in the trajectories is not consistent during the whole interval. Additionally, most of the variance is represented by trajectories with no rotations, on jPCA plane 3. A similar behaviour is found in the projections of the third shuffle control, in fig. B.3(d). The rotations are present, but the direction of the curvature is not consistent throughout the interval. These weak rotations, however, explain slightly more data variance than in the previous shuffle.

Figure B.4 (Cont.) plots the jPCA projections of the original and simulated PMv activity. The trajectories of the original data, in fig. B.4(a), show robust rotations in jPCA planes 1 and 2, and slightly weaker rotations in jPCA plane 3. The loss of rotational structure caused by the shuffle

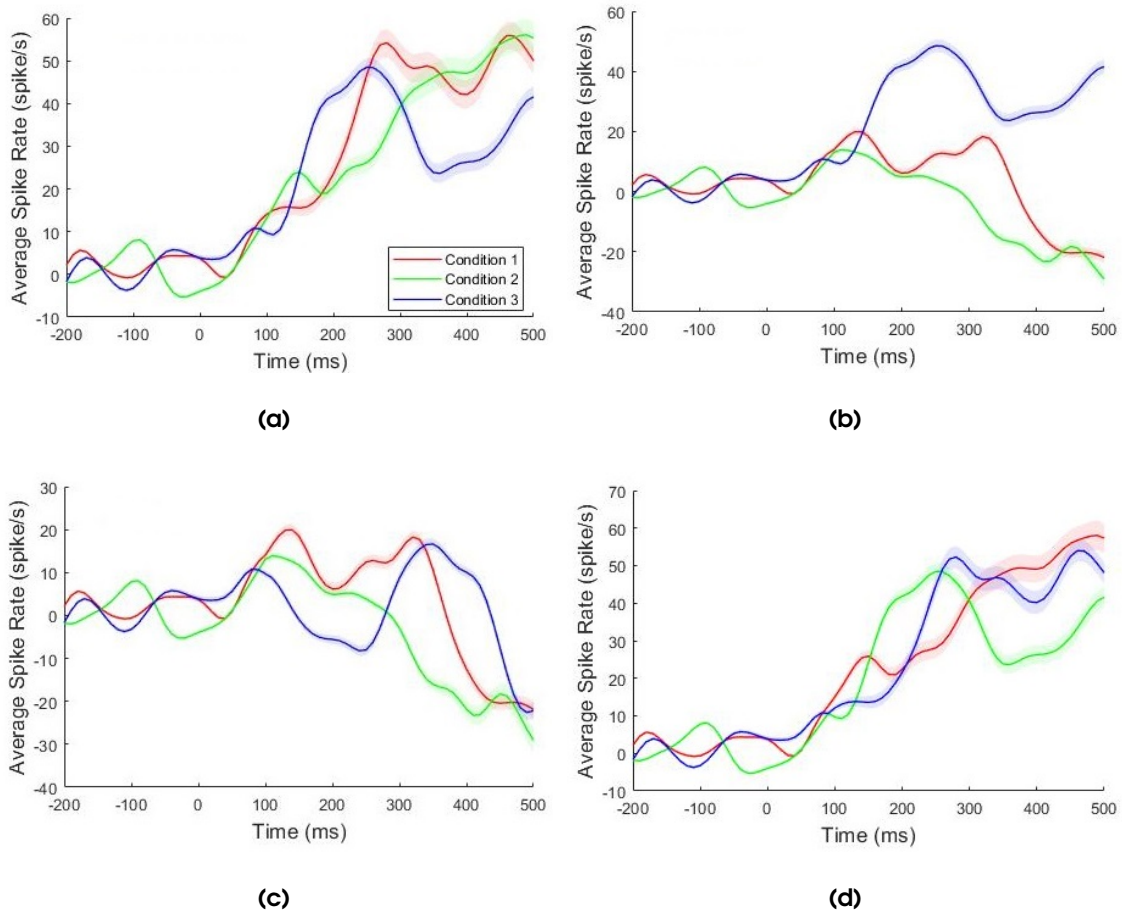

Figure B.1: PSTH of the multi-unit activity recorded in channel 46 of the array implanted in PMd of monkey J3, during 200 ms after object onset, for: **(a)** original data; **(b)** shuffle #1; **(c)** shuffle #2; **(d)** shuffle #3. The standard error is plotted in a lighter shade around the PSTH for each condition.

control 1 is not so obvious in fig. B.4(b). However, looking closely at the rotations on the jPCA planes 2 and 3, the direction of the curvature is not consistent during the whole interval - the angle between the trajectory and its derivative is more likely closer to $-\pi/2$. The shuffle control 2 causes the most of the robustness in plane 2 and to disappear, and the trajectories in plane 3 to practically linearise. Only the first jPCA plane captures strong rotations, although explaining very little data variance. Lastly, the trajectories obtained by the shuffle control 3 have almost linear trajectories on jPCA planes 2 and 3; jPCA plane 1, however, still displays strong rotational structure on the first jPCA plane, capturing little data variance.

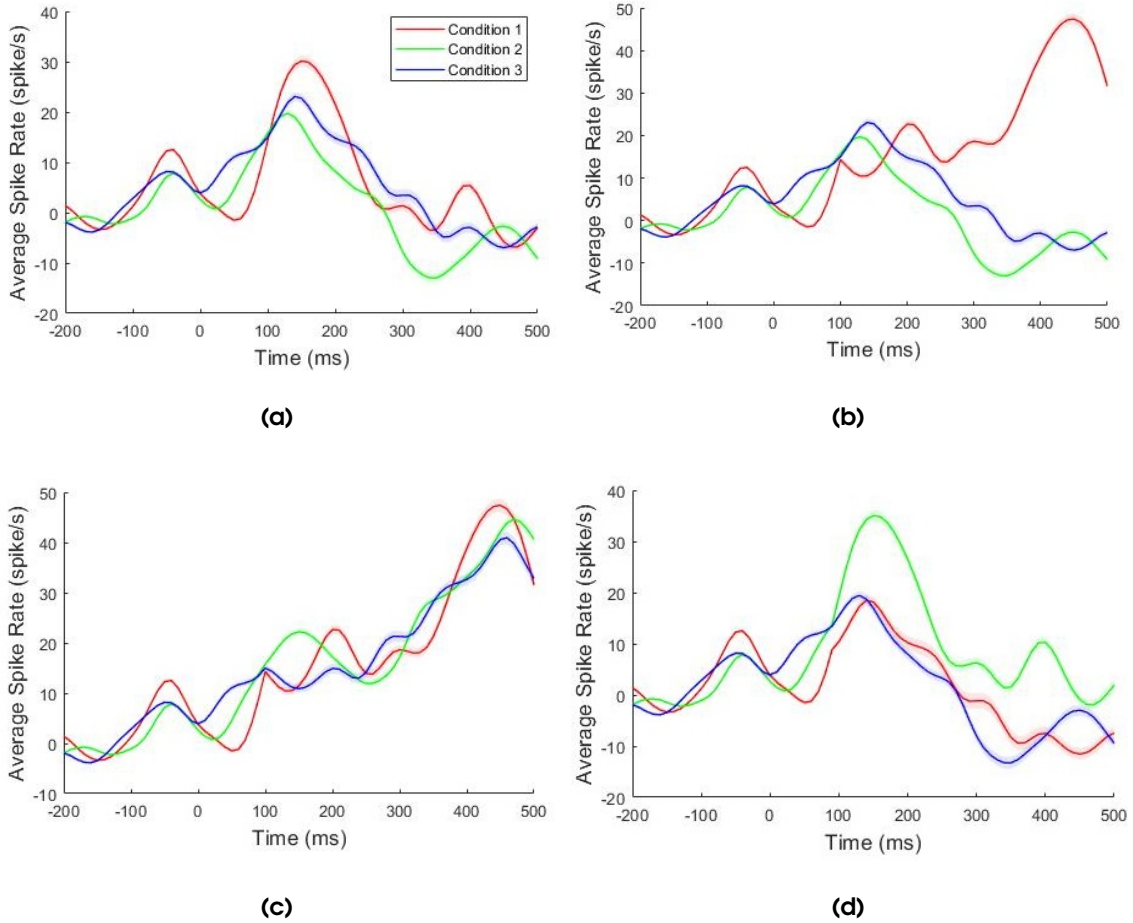

Figure B.2: PSTH of the multi-unit activity recorded in channel 46 of the array implanted in PMv of monkey J3, during 200 ms after object onset, for: **(a)** original data; **(b)** shuffle #1; **(c)** shuffle #2; **(d)** shuffle #3. The standard error is plotted in a lighter shade around the PSTH for each condition.

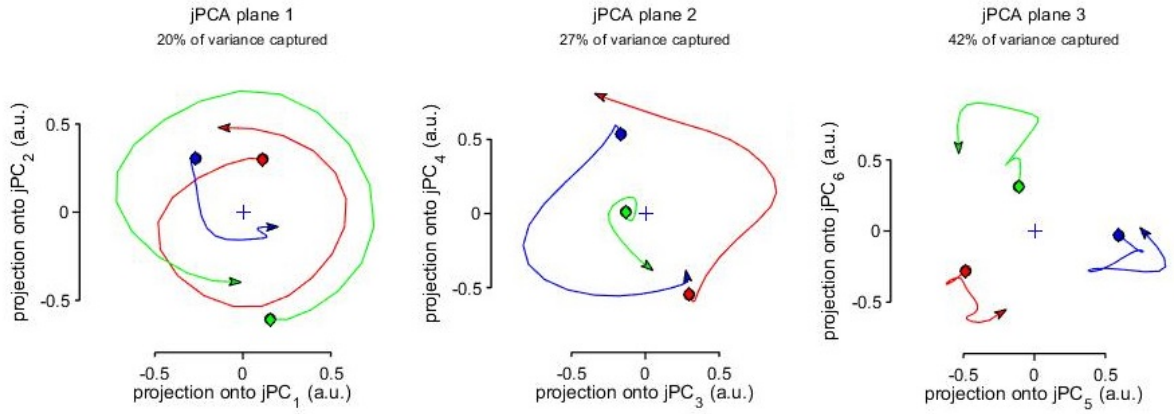

(a)

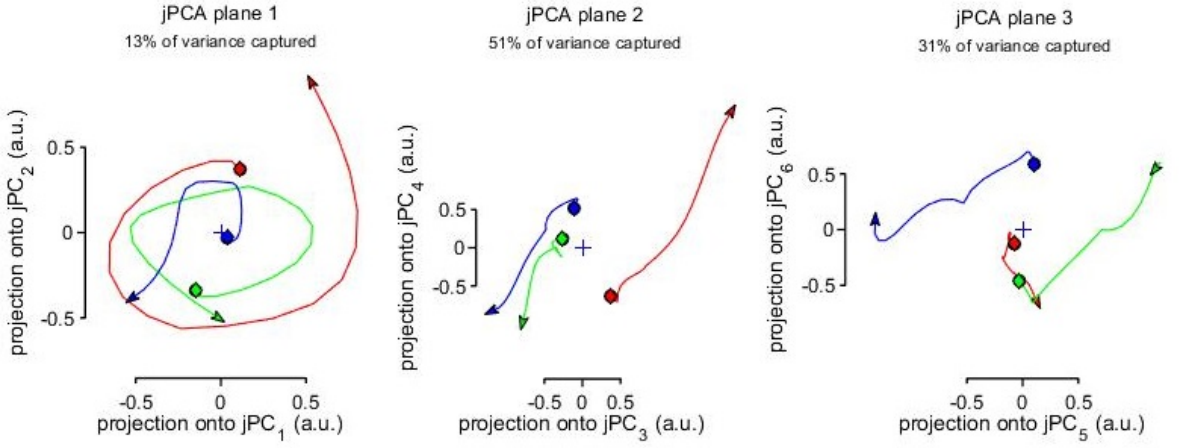

(b)

Figure B.3: State-space projections of the multi-unit activity during the 200 ms after object onset for a PMd population in monkey J3, for: **(a)** original data, **(b)** shuffle #1, **(c)** shuffle #2 and **(d)** shuffle #3. Each trace represents the trajectory for condition 1 (in red), condition 2 (in green) and condition 3 (in blue). The circle and the arrow represent the starting and the end points of the analysis, respectively.

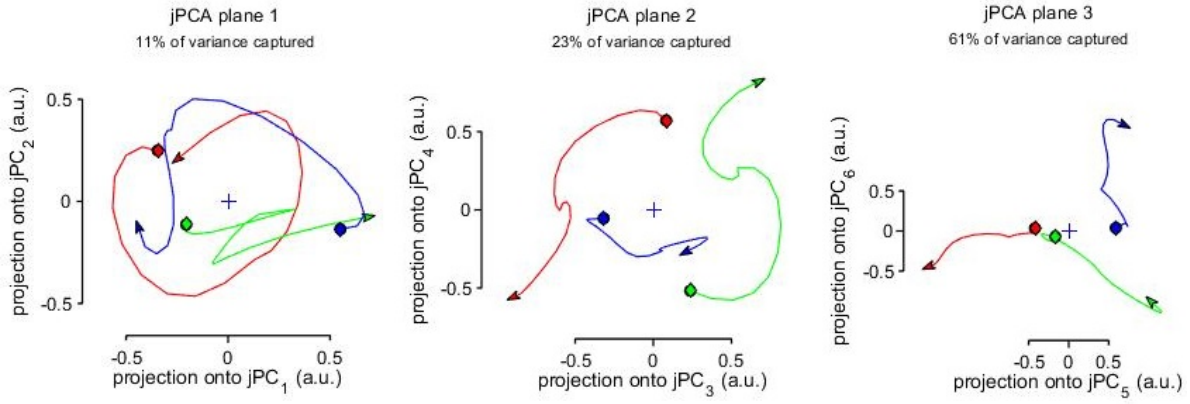

(c)

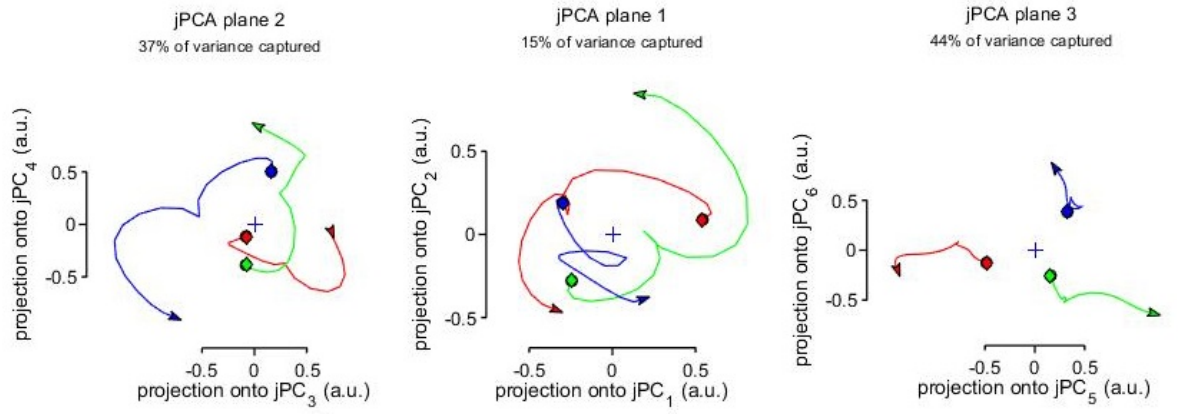

(d)

Figure B.3 (Cont.): State-space projections of the multi-unit activity during the 200 ms after object onset for a PMd population in monkey J3, for: **(a)** original data, **(b)** shuffle #1, **(c)** shuffle #2 and **(d)** shuffle #3. Each trace represents the trajectory for condition 1 (in red), condition 2 (in green) and condition 3 (in blue). The circle and the arrow represent the starting and the end points of the analysis, respectively.

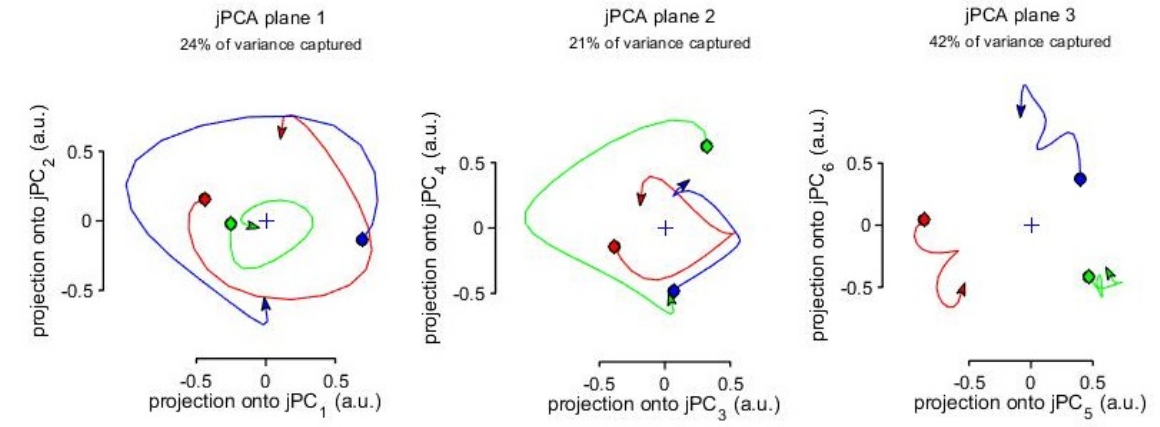

(a)

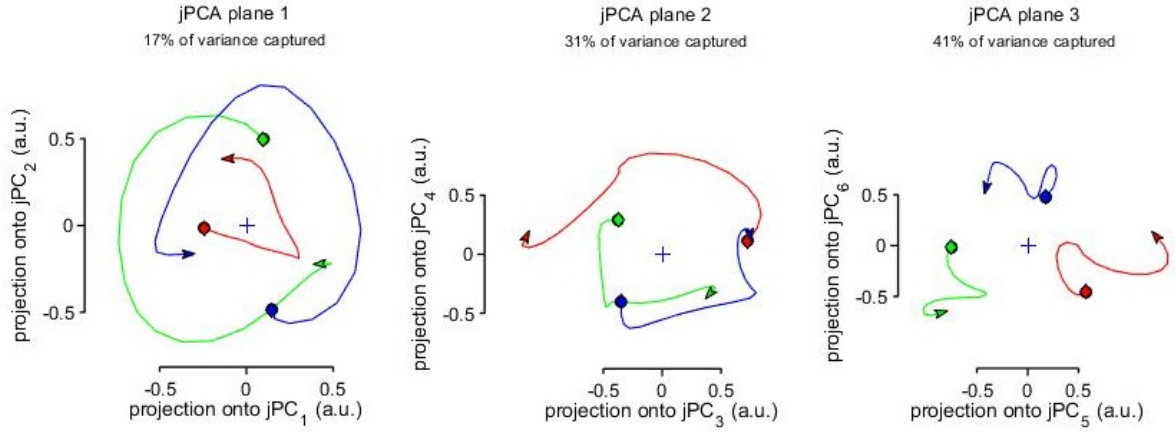

(b)

Figure B.4: State-space projections of the multi-unit activity during the 200 ms after object onset for a PMv population in monkey J3, for: **(a)** original data, **(b)** shuffle #1, **(c)** shuffle #2 and **(d)** shuffle #3. Each trace represents the trajectory for condition 1 (in red), condition 2 (in green) and condition 3 (in blue). The circle and the arrow represent the starting and the end points of the analysis, respectively.

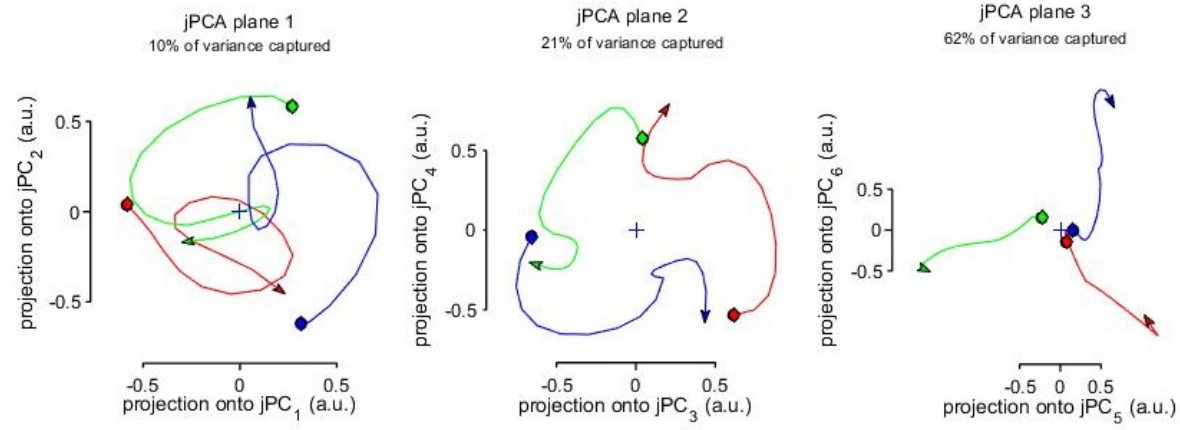

(c)

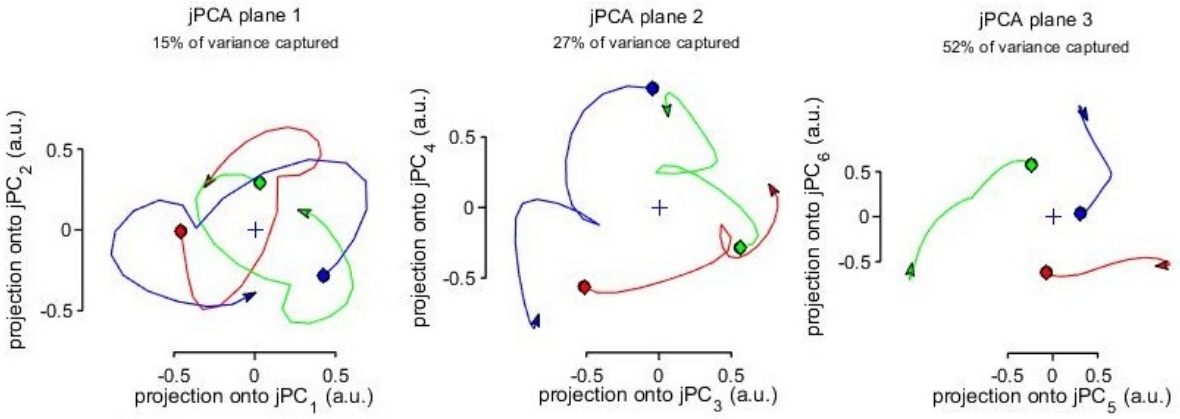

(d)

Figure B.4 (Cont.): State-space projections of the multi-unit activity during the 200 ms after object onset for a PMv population in monkey J3, for: **(a)** original data, **(b)** shuffle #1, **(c)** shuffle #2 and **(d)** shuffle #3. Each trace represents the trajectory for condition 1 (in red), condition 2 (in green) and condition 3 (in blue). The circle and the arrow represent the starting and the end points of the analysis, respectively.