



The
University
Of
Sheffield.

The quantification and characterisation of ramping firing rates and regional differences in the basal ganglia

Adriel Chua Deng Xiang

A thesis submitted in partial fulfilment of the requirements for the degree
of
Doctor of Philosophy

The University of Sheffield
Faculty of Science
Department of Psychology

March, 2023

Acknowledgements

First, I would like to thank my supervisor, Prof. Robert Schmidt, for his guidance and patience throughout my PhD.

I would also like to thank my fellow lab members Alejandro, Lars and Gergo for fruitful work (and social!) discussions.

I am grateful to my parents, Steven and Cecilia, for their unconditional support throughout my university journey. Finally, special thanks to my wonderful significant other, Laura, for her patience and understanding.

Abstract

The basal ganglia are part of the neural circuit involved in inhibitory control. Gradual, consistent changes in firing rate (ramps) have been found in the basal ganglia during 'hold' periods of different behavioural tasks involving inhibitory control. However, the properties of these ramps, such as startpoints, endpoints, amplitude, length and single-trial properties, have not yet been characterised in detail. In order to obtain these properties accurately, we tested and adapted various methods from the wind-ramp and changepoint literature. Specifically, we used pruned exact linear time (PELT) to detect changepoints for ramp detection. Compared to other algorithms, the PELT method was found to have the best performance. Utilising PELT, we then analysed a large data set of electrophysiological recordings in rats performing behavioural tasks to determine the properties of firing rate ramps during hold periods. First, we established that these ramps exist during the hold period in a stop-signal task. Second, we demonstrate that these ramp properties vary widely across the basal ganglia. Third, at the single-trial level, we show that neurons with detected ramps could have different underlying spiking patterns, for example, stepping or ramping. Finally, we investigated the entrainment of neurons in the basal ganglia and found that different waveforms can influence entrainment properties at gamma frequencies.

Contents

1	Introduction	1
1.1	Aim and motivations	1
1.2	Outline	1
2	Background: Basal ganglia, ramps in neuroscience, ramp detection and spike-phase entrainment	4
2.1	Basal ganglia and inhibitory control	4
2.1.1	Inhibitory control and the stop-signal task	5
2.1.2	Architecture and function of the basal ganglia	8
2.1.3	Prototypical and arkypallidal neurons in the GP	10
2.1.4	Models of inhibitory control in the basal ganglia	12
2.2	Ramping firing rate activity	15
2.3	What are ramps useful for	15
2.4	Ramp detection	20
2.4.1	11-sliding window (11sw)	22
2.4.2	Swinging door algorithm (SDA)	24
2.4.3	Changepoint algorithms	27
2.4.4	Opt	28
2.4.5	Pruned Exact Linear Time (PELT)	29
2.5	Spike-phase entrainment	30
3	Quantification of firing rate ramps	34
3.1	Chapter overview	34
3.2	Methods	35
3.2.1	Algorithms	35

3.2.2	Firing rate simulations	36
3.2.3	Quality	37
3.2.4	Evaluation metrics	39
3.3	Results	41
3.3.1	Sensitivity of γ parameter in l1tf	41
3.3.2	PELT penalty hyperparameter is the most robust	42
3.3.3	PELT outperforms other methods in precision of ramp de- tection.	45
3.3.4	PELT shows higher RMSE than l1tf on detected ramps.	45
3.3.5	l1tf is biased towards shorter ramps.	48
3.3.6	Quality measure in simulations	48
3.4	Discussion	53
3.4.1	Comparison with previous approaches	53
4	Properties of ramping activity in the basal ganglia	55
4.1	Introduction	55
4.1.1	Inhibitory control and ramping firing rates.	55
4.2	Methods	56
4.2.1	Dataset	56
4.2.2	Data Analysis	57
4.2.3	PELT algorithm	57
4.3	Results	60
4.3.1	Ramping activity occurs during the hold period throughout the basal ganglia.	60
4.3.2	Ramp properties differ across the basal ganglia.	65
4.3.3	Ramp characteristics do not reflect single trial behaviour	71
4.4	Discussion	77
5	Artifacts in spike-phase entrainment	85
5.1	Chapter overview	85
5.2	Methods	86
5.2.1	Dataset	86
5.2.2	LFP processing	87
5.2.3	Waveform processing	87
5.3	Results	89

5.3.1	Waveform differences between arkypallidal and prototypical neurons	89
5.3.2	Differences in waveforms result in different entrainment phases.	89
5.4	Discussion	92
6	General discussion	95
6.1	Overview	95
	Appendices	112
A	Chapter 3: Quantification of firing rate ramps	113
B	Chapter 4: Properties of ramping activity in the basal ganglia	117
B.1	Detected ramps aligned to Nose in	117
B.2	Firing rates with no detected ramps aligned to Nose in	123
B.3	Detected ramps aligned to Go cue	128
B.4	Firing rates with no detected ramps aligned to Go cue	133
B.5	Single trial analysis of positive detected ramps.	138
B.6	Single trial analysis of negative detected ramps.	144

Chapter 1

Introduction

1.1 Aim and motivations

Firing rate ramps have been identified across regions of the Basal Ganglia. These have mostly been qualitatively assessed and categorised; however little progress has been made on quantitatively characterising ramp features. In this thesis, we, therefore, demonstrate how to better characterise ramps, comparing several algorithms. Using this characterisation, we then compare ramps across regions of the basal ganglia, which have not previously been extensively compared. Another debated issue is whether these ramps are reflected in single trials or are made up of many steps in single trials. Thus we investigate whether average firing ramp rates are reflected in single trial data. Finally, we explored spike-phase entrainment in gamma oscillations and GP neurons and looked at artifacts in spike-phase entrainment.

An overview of the contributions of each chapter are listed below:

1.2 Outline

Chapter 2: Background: Basal ganglia, ramps in neuroscience, ramp detection and spike-phase entrainment

In this chapter, we give an overview of background information relevant to the thesis. First, we discuss the concept of inhibitory control and how this is mea-

sured through different behavioural tasks, such as the stop-signal task. We then introduce firing rate ramps typically studied during the hold period of these tasks. Next, we highlight methods of ramp detection in other fields, such as wind power, which are typically more accurately characterised than currently done in neuroscience. We then discuss artifacts in spike phase entrainment in neuroscience.

Chapter 3: Quantification of firing rate ramps

In Chapter 3, we investigate different algorithms that can be used to detect firing rate ramps with more accuracy than previously done in neuroscience. First, the algorithms introduced in the introduction are tested and compared on simulated firing rate ramps. We then quantified and compared the sensitivity of hyperparameters in various algorithms and developed the PELT algorithm to be suitable for detecting ramps.

Contributions:

- Developed a ramp detection method based on the PELT algorithm.
- Quantified the sensitivity of hyperparameters in l1tf and PELT.
- PELT algorithm outperforms other methods in precision of ramp detection in our simulations.

Chapter 4: Properties of ramping activity in the basal ganglia

In this chapter, we apply the ramping characterisation method to datasets taken from a stop-signal task with recordings in the basal ganglia. We then investigated the single trial behaviour of trial-averaged ramping neurons.

Contributions:

- Established the existence of ramps in the basal ganglia during the hold period.
- Found heterogenous ramp properties across various regions in the basal ganglia.
- Developed method to identify ramping vs stepping in single trial data.
- Found that neurons exhibiting ramping behaviour in their average firing rate consisted of both ramping and stepping patterns at the single-trial level.

Publications in preparation:

- **Chua, A. D. X.** & Schmidt, R. (2022). Heterogeneous properties of ramping activity in the basal ganglia. *In preparation*

Chapter 5: Artifacts in spike-phase entrainment

In chapter 5, we examine spike-phase entrainment in arkypallidal and prototypical neurons. We then looked at the waveform differences between putative arkypallidal and prototypical neurons. Finally, we quantified how these waveforms can influence spike-phase entrainment.

Contributions:

- Found waveform differences in putative arkypallidal and prototypical neurons.
- Demonstrate how different waveforms can affect spike-phase entrainment.

Chapter 6: General discussion

Here we discuss the thesis contents in its broader context. We begin by summarising the contents of each chapter and then describe how the work could affect broader research in neuroscience.

Chapter 2

Background: Basal ganglia, ramps in neuroscience, ramp detection and spike-phase entrainment

2.1 Basal ganglia and inhibitory control

In this chapter, we introduce the context of the thesis. We first discuss the relevant definitions of inhibitory control and tasks used to study it, such as the stop-signal task. The datasets used for each of the thesis chapters also utilise this task. We then discuss regions of the brain implicated in inhibitory control, particularly the architecture, function and models of the basal ganglia. The thesis has two main focuses, first, relating to quantification and analysis of ramping in neural signals and second, understanding artifacts in spike phase entrainment. We, therefore, introduce previous experimental studies identifying ramping of neural activity in the basal ganglia during inhibitory control tasks, followed by a review of algorithms that have been developed more generally in signal processing domains to identify ramps in temporal signals. Finally, we discuss spike phase entrainment of neural signals relevant to the final results chapter 5.

2.1.1 Inhibitory control and the stop-signal task

Inhibitory control is a core executive function that allows the inhibition of undesirable or even dangerous actions and thoughts. Inhibitory control is important in executing daily tasks because it helps us to respond accordingly to stimuli (Bari and Robbins, 2013; Diamond, 2013; Logan, 1994). For example, when you see the green man light up while waiting at a crossing, your natural reaction would be to walk across the road. However, if you notice a car from a distance not slowing down, you would have to quickly stop yourself or move back to the pavement to avoid getting hit.

Different forms of inhibitory control exist, such as reactive and proactive inhibition of behavioural responses (Schmidt et al., 2013; Gu and Berke, 2022). Reactive inhibition is where you need to inhibit and cancel an upcoming action based on new external stimuli, for example, seeing the car in the previous example. Proactive inhibition is akin to a preparatory phase, where you are more prepared to stop if needed; there is proactive suppression of the response. Two subtypes of proactive inhibition are action restraint (withholding an action) and postponing (waiting to initiate the action) (see Bari and Robbins (2013) for a review of impulsivity types).

A lack of response inhibition, also known as impulsivity, can be defined as the difficulty of inhibiting one's own thoughts and behavioural actions. Impulsivity has been linked to various disorders (Robbins et al., 2012; Chambers et al., 2009), such as addiction (Crews and Boettiger, 2009; Stevens et al., 2014), attention deficit hyperactivity disorder (ADHD) (Nigg, 2001; Winstanley et al., 2006), mania and schizophrenia (Gut-Fayand et al., 2001).

In order to understand the possible neural mechanisms behind inhibitory control, many types of behavioural tasks have been developed. Response inhibition is commonly studied with the stop-signal task (SST) (Schmidt et al., 2013; Mallet et al., 2016; Mirzaei et al., 2017; Eagle et al., 2011; Lipszyc and Schachar, 2010) (see also the review of the SST by Verbruggen and Logan (2009)). These kinds of behavioural tasks often include hold periods, during which the animal or participant must exert action restraint and wait (e.g. for a Go cue) to prevent premature responses.

The stop-signal task has been utilised in many studies to investigate action inhibition (Logan (1994); Verbruggen and Logan (2009), see fig. 2.1). When the task

begins, the animal or participant is expected to prepare for a Go stimulus. The time between the cue for the task start and the Go stimulus is called the 'hold' period. Once the Go stimulus is presented, an immediate response is expected, usually a specific movement to be carried out. This Go signal could be in the form of a green light, for example. On a subset of these trials, the 'Go' stimulus would be followed by a stop-signal, where the subject is signalled to halt action initiation. The number of stop trials is kept to a small fraction of the total trials, and the occurrence of Go and stop trials is randomised to ensure the subject does not expect them. Some experiments adjust the time between the Go cue and the stop-signal such that subjects have equal numbers of failures and successes in the task, while others use a fixed delay (Verbruggen and Logan, 2009).

When rats are the task subjects, they typically start with their Nose in a 'home' nose-poke port. Presentation of the Go stimulus indicates that they must leave the home centre port and move to another port. Correct responses to the stimulus results in a reward. In stop trials, the stop-signal (a tone or light, for example) is presented shortly after the Go stimulus according to some delay. The time between the Go stimulus and the stop cue is called the stop-signal delay (SSD). A correct cancelling of the action initiation (remaining in the home port for a short period) results in a reward. There are also several variations of this task; for example, the Go cue might indicate to move to one of the multiple ports depending on the pitch of an auditory tone (Leventhal et al., 2012).

In the stop-signal task, information on two important processes in inhibitory control can be gathered. First, information about the subject's reaction to the Go stimulus can be accessed via both the Go cue reaction time (time from the Go cue to action, such as moving out of the nose port) and the accuracy of response to the cue. Second, the stop-signal reaction time (SSRT) can be used to evaluate action suppression; it is based on the period between the stop-signal and the time the action is cancelled. It tells us how quickly the subject can cancel the action which may have been initiated following the Go cue.

The stop-signal task has been used to study inhibition in many domains, including cognitive and developmental psychology and in clinical disorders (see Verbruggen and Logan (2009) for a review). The SSRT in the stop-signal task has been used to investigate inhibitory control in various disorders, including obsessive-compulsive disorder (Mar et al., 2022; Lipszyc and Schachar, 2010), and ADHD (Lipszyc and

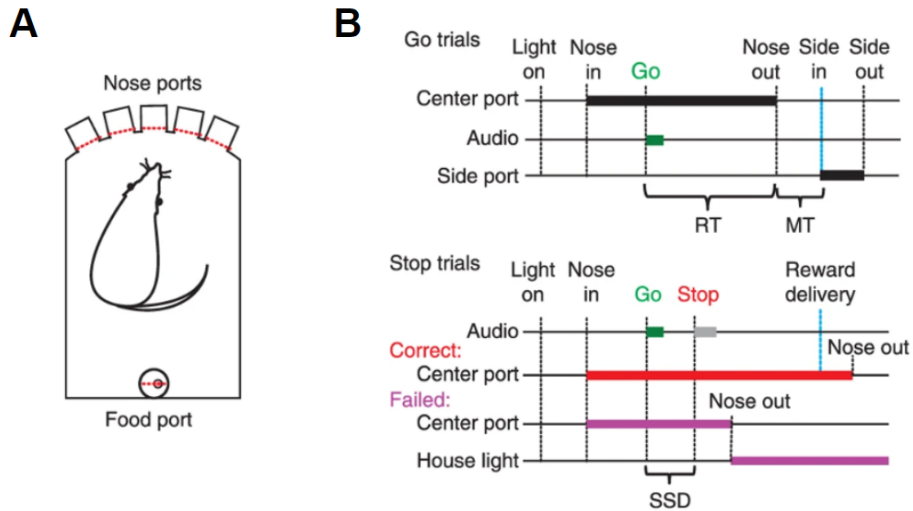


Figure 2.1: Stop-signal task A. Experimental setup of rat in SST. At the start of the task, the rat places his nose into the centre nose-poke port. The rat remains in the centre port until the Go cue. B. Schematic of trial progression for both Go trials (top) and Stop trials (bottom). Left labels refer to rat position in the setup and activation of the audio tone. Top labels are signals for the start of the trial (light on), the Go and stop cues and the position of the rat's nose. RT refers to the reaction time between the Go cue and movement; MT is the movement time between the centre port and the side ports; SSD is the stop-signal delay between the Go cue and stop-signal. Thicker lines denote where cues are active and which port the rat is currently inside. In the 'Go' trials, the rat's nose is in the centre port (black line) until the occurrence of the Go cue (audio cue). The rats then initiate a movement to the side port and can then collect a reward from the food port for a successful action. In stop trials, a Stop cue is given shortly after the Go cue according to the SSD. Where the correct response is completed (cancel the action and do not move), the animal should stay in the centre port (red line) until rewarded. In failed stop trials (the action is initiated, and the rat removes their nose from the centre port), no reward is given. Image adapted from Schmidt et al. (2013).

Schachar, 2010). For example, ADHD-diagnosed participants were found to have slower SSRTs while completing the SSRT task (Lijffijt et al., 2005). Understanding the neural mechanisms behind inhibitory processes has implications for understanding behaviour and treating disorders related to deficits in inhibitory processes.

The neural mechanisms involved in inhibitory control have been studied using the stop-signal task. These have shaped our understanding, highlighting the basal ganglia as a key region. For example, Mallet et al. (2016) found that a subpopulation of cells (arkypallidal) are responsible for the stop-signal within the Striatum. Computational models have also been developed in relation to basal ganglia processes during the stop-signal task (Schmidt et al., 2013). Inhibitory control and computational models concerning the basal ganglia are discussed further in the following sections.

2.1.2 Architecture and function of the basal ganglia

The basal ganglia are a set of subcortical nuclei which have widely been implicated in inhibitory control. Inputs to the basal ganglia originate mostly from cortical regions; however, some also come from the thalamus (Lanciego et al., 2012). Voluntary motor movement signals, for example, originate in the premotor and motor areas. The basal ganglia regulate these behaviours, ensuring the selection of appropriate movements.

One characterisation of how this is achieved is the 'rate model' of the basal ganglia, which was based on anatomical nuclei, their connections and neurochemistry (Alexander and Crutcher, 1990; Nelson and Kreitzer, 2014). It should be noted that other viewpoints and models of the basal ganglia exist. The rate model suggests that ensuring appropriate action selection is primarily achieved through the basal ganglia tonic inhibition on the thalamus. When an action is to be initiated, the basal ganglia's output nuclei releases its inhibition. This output region is known as the globus pallidus internal (GPi) in humans, and the substantia nigra pars reticulata (SNr) in rats.

The rate model highlights two pathways in the striatum of the basal ganglia (Alexander and Crutcher, 1990), the 'direct', which is associated with movement initiation and the 'indirect', associated with movement inhibition. The pathways

are shown in relation to the basal ganglia network architecture in fig. 2.2. Each of these pathways involves one of two subpopulations of medium spiny neurons (MSN) in the striatum, which can be segregated based on their distinct functional properties (Gerfen and Surmeier, 2011). The 'direct' pathway involves GABAergic inhibition of the SNr/GPi through MSNs with D₁ dopamine receptors (see fig. 2.2, green lines). The 'indirect' pathway has MSNs with D₂ type dopamine receptors, providing GABAergic inhibition on the GPe, which in turn inhibits the output structure (GPi/SNr) and also the STN (fig. 2.2, blue lines).

The model was supported by early computational modelling of the basal ganglia network in action selection (Gurney et al., 2001a,b). Furthermore, experimental evidence also supports the two pathways. For example, optogenetic D₁ activation increases motor movement output, whereas activating D₂ MSNs leads to a decrease (Kravitz et al., 2010). Although the rate model explained many aspects of basal ganglia network activity and the resulting behaviours, the model is unable to explain all experimental findings, particularly more recent works that have questioned some of the model assumptions.

First, it has been found that at movement onset, some pallidal cells of the output GPi and thalamic cells have similar firing patterns, questioning the effect of GPi inhibition (Goldberg et al., 2013; Schwab et al., 2020). Second, GPi lesions have been found not to affect thalamic output (Goldberg and Fee, 2012; Inase et al., 1996; Goldberg et al., 2013), a finding at odds with the GPi as a major input region affecting thalamic firing. Furthermore, two opposing disorders, one with hypokinetic (PD) and the other with hyperkinetic behaviour (dystonia), can both be treated by lesioning and simulation of the GPi (Nelson and Kreitzer, 2014; Magnusson and Leventhal, 2021). In contrast, different approaches would be expected to be needed, targeting the GPi to produce opposite effects- either suppress unwanted actions or enable action initiation. Finally, the distinction of two pathways in the rate model does not explain why both D₁ and D₂ MSNs can be coactive (Calabresi et al., 2014). Magnusson and Leventhal (2021) reviewed the evidence consistent and conflicting this 'rate model' and suggested that many of these paradoxes can be explained by thalamocortical loops being the fundamental drivers of motor output.

A third additional 'hyperdirect' pathway between the STN and GPi (fig. 2.2, red lines) has also been considered, which is faster than both the 'direct' and 'indi-

rect' pathways of the rate model, in that response time to stop-signals tends to be quicker. As a result, it has been hypothesised that it is responsible for the fast stopping of behaviours (Nambu et al., 2002). Excitation of the STN leads to increased excitation of the GPi, increasing inhibition and stopping action selection. Further evidence of the STN's role in inhibitory control is from treatments of Parkinson's Disease. Deep brain stimulation (DBS) of the STN was found to alleviate symptoms (Lachenmayer et al., 2021). Although the exact mechanisms of how DBS works are still unclear and debated (Alhourani et al., 2015), this may suggest that increasing STN activity may inhibit unwanted movements. Nevertheless, recent studies utilising optogenetic methods further support the role of the STN in action selection; for example, a brief excitation of STN cells in mice was found to pause or delay licking behaviours (Fife et al., 2017).

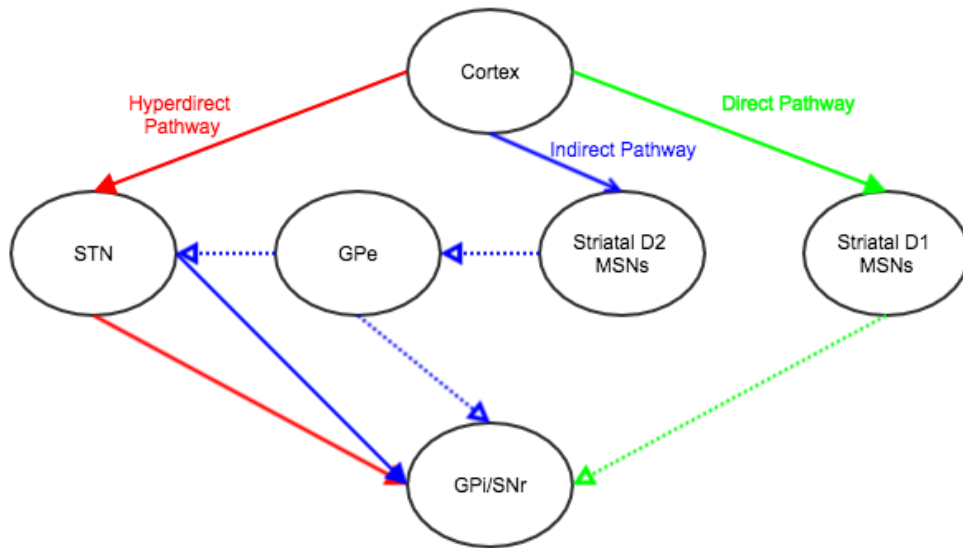


Figure 2.2: Classical View of the 'direct' (green), 'indirect' (blue) and 'hyperdirect' (red) basal ganglia pathways. Dotted lines with empty arrowheads represent inhibitory connections. Solid lines with filled arrowheads represent excitatory connections. Image taken from Chua (2017).

2.1.3 Prototypical and arkypallidal neurons in the GP

The external globus pallidus (GPe) has been implicated in inhibitory control during action selection for several reasons. First, the GPe has widespread inhibitory

connections to most basal ganglia nuclei, including the STN, GPi and the Str (Lanciego et al., 2012; Mallet et al., 2016). Second, GPe dysfunction and pathological activity has been found in several motor disorders, including PD, Huntington's Disease (HD) and dystonia (Nambu et al., 2011). In PD, exaggerated oscillatory activity has been found in the GPe-STN network (Mallet et al., 2008; Brown et al., 2001; Chan et al., 2011). Successful treatments to alleviate symptoms of HD have involved electrical stimulation of GPe neurons, further suggesting a role in the disease (Gonzalez et al., 2014; Beste and Saft, 2015; Nagel et al., 2015), either directly or through network interactions.

Mallet et al. (2012) suggested that the GPe consists of two distinct populations of neurons, namely the arky pallidal (arky) and prototypical (proto) cells. These subpopulations have different electrophysiological properties and axonal connectivity. Therefore, their inclusion in models may lead to different basal ganglia dynamics than predicted by earlier models formulated before arky and protos were identified. Studying these subpopulations may also help us better understand inhibitory control in the basal ganglia (Schmidt and Berke, 2017).

Protos constitute the majority of cells, at around 60-70% of those identified (Dodson et al., 2015; Mallet et al., 2012; Dong et al., 2021). They are distinctive in that they have both high and consistent firing rates (Abdi et al., 2015; Dodson et al., 2015). Further subtypes can be identified within the GP populations based on molecular expression (Dong et al., 2021). Protos that express the calcium-binding protein parvalbumin (PV) constitute about 55% of the neurons in the GP, although some protos do not express PV. Prototypical neurons receive inputs from the striatum, STN, cortex and thalamus (Milardi et al., 2015; Yasukawa et al., 2004). Although they primarily receive GABAergic inputs from the striatal D2 MSNs and glutaminergic inputs from the STN. Protos main output is directed towards the STN and SNr, but they also innervate the striatum and local neurons (Nambu and Llinas, 1997; Kita and Kita, 1994; Mallet et al., 2012; Hegeman et al., 2016).

Arkys consists of about 25% of all GP neurons, and unlike protos, they tend to have irregular and lower firing rates (Abdi et al., 2015; Dodson et al., 2015). Instead of PV, arkys express preproenkephalin and the transcription factors *Foxp2* and *npas1+* (Mallet et al., 2012). Striatal D2 MSNs also innervate arkys. This input seems to be significantly weaker to arkys than to the protos (Dodson et al., 2015; Glajch et al., 2016; Chuhma et al., 2011). The main output of arkys is directed

towards the Str, but they also have local collaterals to other nuclei within the GP (Nambu and Llinas, 1997; Kita and Kita, 1994; Mallet et al., 2012; Hegeman et al., 2016). These arky cells fire in phase with STN neurons and fire anti-phase with GPe proto neurons (Mallet et al., 2012). Protos tend to receive stronger inputs from the STN than the arkys (Aristieta and Gittis, 2021).

In vivo optogenetic activation of striatal D2 MSN was found to lead to time-locked inhibition of prototypical GPe, and increase of arky neurons activity (Aristieta and Gittis, 2021). Protos receive more striatal D2 MSN input compared to arky. When the STN excitatory inputs to the GP were optogenetically inhibited, proto activity was decreased, time-locked with the STN inhibition, while arky activity increased. When these STN inputs were instead optogenetically excited, the activity of protos increased, time-locked with the STN excitation. Meanwhile, arky activity was found to decrease. When both the D2 striatal MSNs and STN neurons were simultaneously stimulated, arkys responded stronger than when either was activated alone, suggesting a more complicated interaction between them. In general, protos were found to inhibit arkys and protos while arkys outputs to prototypic neurons are much smaller in size.

In the stop-signal task, the arky cells of the GP have been found to be more selective to stop signals rather than Go signals. They likely have a greater role than the proto sub-population, with responses around 50ms, compared to much slower and later peak proto response (Mallet et al., 2016). Arkys may suppress the activity of the striatum during the stop-signal, as evidenced by their quicker response timing and connectivity. Optogenetic studies where activation of *npas1*-expressing neurons in the GPe also support the arky neuron's role in action suppression, as activation of these has been shown to suppress motor output (Glajch et al., 2016).

2.1.4 Models of inhibitory control in the basal ganglia

Numerous models have been proposed to explain how the basal ganglia network is involved in inhibitory control.

Schmidt et al. (2013) developed a race model of inhibitory control in the basal ganglia based on earlier experimental work. This data was from the stop-signal task in rats, whilst single unit electrophysiology was recorded from the basal ganglia (see earlier fig. 2.1A,B for SST setup). The SNr only responded to correct stop

trials, while the STN responded to both correct stop and failed stop trials. This implies that the STN might only successfully transmit a signal to the SNr on correct stop trials. On failed stop trials, however, it is suggested that the go signal from the striatal neurons arrives at the SNr faster, thus initiating action. However, when the STN stop-signal arrives at the SNr faster than the striatal go signal, the animal is able to stop successfully, and there is no subsequent action initiation (see fig. 2.3). This suggests that there is a 'race' in which the direct pathway generates the go signal and the indirect pathway generates the stop-signal. When comparing reaction time lengths between rats, those with longer reaction times were generally more successful in responding to the stop cue signal and stopping the action compared to rats with shorter reaction times. However, it is unlikely that the transient activity between STN-SNr completely inhibits the SNr, as this signal lasts only around 15ms. Therefore, successful responding to the stop-signal likely requires striatal activity to be reduced in addition.

While many models of inhibitory control in the basal ganglia have been developed, few computational models have included both arky and proto in their model. For example, earlier models such as Nevado-Holgado et al. (2014) and Bogacz (2015) did not include the arky-striatal pathway. The distinction of arky-proto cells in the model of Wei and Wang (2016) found that the GPe-striatum connection is crucial for inhibitory control.

The neurophysiology of the basal ganglia, particularly the processes and connections of the GP, and their relationship to stopping behaviours have been encapsulated by the 'pause-then-Cancel' model (Schmidt and Berke, 2017). The model is based on experimental evidence from Schmidt et al. (2013) discussed above and Mallet et al. (2016). In Mallet et al. (2016), enhanced activity of arky cells was found when stop cues were presented. On correct stop trials, the timing of this cue was around the time when it could be the driver of Str suppression.

The 'pause-then-Cancel' model highlights that action suppression is served via two different mechanisms. First, the STN enables quick suppression of the basal ganglia output by sending diffuse excitation directly to the GPi. After a short delay, the second mechanism involves arky cells inhibiting striatal MSNs, leading to a reduction in action initiation and a resulting stopping effect on the basal ganglia output.




Figure 2.3: Race model of Schmidt et al. (2013). Race model schematic relating to basal ganglia regions. The SNr receives input from both the striatum (STR) and STN. SNr tonically inhibits the motor outputs. When this inhibition is released, actions can be initiated. Inhibition of the SNr by the STN forms the 'Go' signal, whereas excitation from the STN gives the 'Stop' signal. Image adapted from Schmidt et al. (2013).

Although the model was formulated from experimental results on rats, recently, the model has been investigated with human participants and suggests that the two stages (pause and cancel) can map onto human experimental findings (Diesburg and Wessel, 2021).

Recent experimental data have identified connections from arky and GPe neurons, which project to cortex (cortico-pallido-cortical loop); these pathways and connections were added to a previously developed spiking model of different basal ganglia nuclei (Goenner et al., 2021), with neuron responses fit to the experimental data. The model was able to reproduce the neuronal and behavioural activity of a simulated stop-signal task. They suggested several processes may complement each other to stop a behaviour, including the identified cortical input to arky neurons and the GPe to cortex projections. This process may work alongside the shorter inhibition of GPe-arky on striatal activity.

2.2 Ramping firing rate activity

As discussed earlier, the stop-signal reaction time task has a 'hold' period, where the subject must wait after the trial onset until the Go cue is presented. It has been suggested that actions may be inhibited during this period by ramping firing rates (London et al., 2018; Emmons et al., 2017; Mirzaei et al., 2017; Donnelly et al., 2015; Narayanan, 2016). Firing rate ramps occur when spiking activity increases or decreases over an extended time period, and the shape of these ramps can be consistent or non-linear (Zhao and Kording, 2018). Ramps have been found in various brain regions, such as the striatum (London et al., 2018; Emmons et al., 2017; Mirzaei et al., 2017; Donnelly et al., 2015), the lateral intraparietal area (LIP) (Roitman and Shadlen, 2002; Meister et al., 2013; Latimer et al., 2015) and other cortical regions (Narayanan, 2016). Firing rate ramps can vary in length from as short as two seconds up to 15 seconds (Emmons et al., 2017; London et al., 2018). The varying statistical properties of ramps in different tasks and regions, and the existence of non-linear ramps, make their detection and characterisation challenging.

2.3 What are ramps useful for

Ramping activity is able to capture the accumulation of temporal evidence, this means that as the expectation of an event occurring increases due to incoming evidence, neuronal activity can also increase (or decrease). This has been captured by drift-diffusion models and integrative models in past works (Latimer et al., 2015; Simen et al., 2011). In the drift-diffusion models, the slope of the drift is affected by sensory evidence, and in each trial, activity increases(or decreases) until an upper(or lower) threshold is hit.

A "hazard" function can mathematically represent temporal expectation (Narayanan, 2016). For example, if an event that is predicted to happen in a certain period of time doesn't occur, the likelihood it will happen after that time then increases. While making decisions, organisms utilize temporal information since it becomes increasingly apparent when events will occur as time passes. If the neuron's activity is a step instead of a ramp, certainty regarding events cannot be increased as steps do not have a varying maximum amplitude to encode timing. Stepping, especially in single trials is unable to represent varying degrees of confidence over time.

Another simple function of ramps is to relate their amplitude directly to timing. For example, during hold periods where animals are expected to 'hold' movement between 5-10 seconds, the amplitude of the firing rate could simply increase over time. The amplitude could be used as an estimate of how much time the animal has been holding.

The function of neural ramping activity has been hypothesized to help in temporal regulation, for example, accurately controlling our actions in time for decision making and basic survival (Narayanan, 2016). Ramping may be a neural way to implement evidence accumulation, the accumulation of incoming sensory information over time until enough has been gathered to reach a decision threshold and make a decision.

Many behavioural tasks include some form of temporal expectation. For example, in the stop-signal task, when the animal is trained, it should be able to expect the Go cue to arrive within the time it was typically trained (for example, one second). The longer the Go cue does not come on, the more likely the Go cue would then appear. Thus, the probability of an event increases, as the event fails to occur over time. To accurately perform on these tasks, animals should utilize this temporal information from past trials.

However, ramping activity might serve many different functions depending on the behavioural task and has, besides evidence accumulation, also been connected to the temporal coordination of movements to achieve behavioural goals (temporal control) (Narayanan, 2016), action restraint (Frank, 2006; London et al., 2018; Emmons et al., 2017; Mirzaei et al., 2017; Donnelly et al., 2015), reward prediction, and time estimation.

Ramping may be related to temporal control in the ventral Str (VS). Donnelly et al. (2015) investigated ramping during premature responses to stimuli in rats, where the response is too early and occurs before the target stimuli are presented and the appropriate action should be initiated. Impulsivity may be involved in premature actions (where an unplanned quick decision or action occurs). This may relate to prefrontal control of behavioural control and timing/ temporal control of the basal ganglia. In the study, a five-choice serial reaction time test (5-CSRTT) task was used. Rats performed the task in a chamber with five different nose-poke ports. On each trial, there was first a fixed delay of 5s; after this, a light stimulus

was activated in one of the five ports, signalling the rat to nose-poke that port. If the rat completed the poke within 5s, they were rewarded with a food pellet. A premature response, with the rat nose-poking before the initial delay or an incorrect response (wrong port) was punished by a timeout period of 5s. Both positive and negative ramping activity was found in recorded regions of the medial prefrontal cortex (mPFC), and in the VS. In the VS, 50.4% of cells showed ramping. Negative ramps were found to begin earlier on premature trials but were similar to correct trials in other characteristics, such as the gradient. In positive trials, it was not as apparent, but the pre-ramp min and post-ramp maximum were reached. The ramps tended to increase/decrease at similar rates (gradients of slope) and reach a similar maximum and minimum rate, regardless of the delay. In the premature trials ramping was found to begin earlier, however. They suggest that the premature responses may be due to the rats failing to correctly time the initiation of the waiting process. They suggest that ramping may be an internal time representation, which is too early on premature responses. In premature trials, ramping activity tended to start earlier but reached similar maxima and minima during the responses. The ramp did not have a steeper gradient but seemed to start earlier. Overall, the premature responses were not associated with a lower peak of ramping activity or a steeper gradient. When a variable delay period was implemented, ramps occurred until the earliest possible time and remained until the nose poke. It is possible premature responses may not be a waiting failure but incorrect timing of the stimulus. Donnelly et al. (2015) suggested that this is likely initiated in structures outside of the VS and mPFC.

In the dorsal Str, responses can be excitatory and inhibitory towards food-related actions, for example, approach to food and consumption. London et al. (2018) investigated this in a simple task with rats. On each trial, rats were trained to remain in a 'trigger zone' area for six seconds until a tone sounded when they had six seconds to move towards the pellet dispenser. They were rewarded with a food pellet if they reached the pellet dispenser within the set time. In around 40% of the recorded neurons, ramping changes were found. These ramps preceded the initiation of a trial rather than at consumption time, supporting the idea of ramping being involved in upcoming action expectations or planning future actions. Ramps have been found to end when an action is initiated, suggesting that the ramps may reflect upcoming actions. Differences were identified between

population responses measured with calcium imaging and at spike level with electrophysiology. At spike level, 2/3 were found to be increasing, and 1/3 were decreasing. However, at the population level, the ramps appeared to be positive. They suggested this might reflect more positive ramping at the individual level or due to the recording method.

In an interval timing task, Emmons et al. (2017) found ramping activity across frontostriatal circuits, prominent in dorsal medial striatal (DMS) neurons. Animals were trained to complete a 12s fixed-interval task. At the start of the task, a houselight is presented, and the rat must wait for 12s to respond by pressing a lever to receive a water reward. A click sound was then played, and the houselight turned off to signal the end of the trial. Random intertrial intervals were used between 6 and 12s. Ramping activity in frontostriatal circuits might be related to interval timing and temporal signalling. Frontal and striatal regions share extensive neuronal connections, and activity in ramps is correlated between the Str and medial frontal cortex (MFC). When the MFC is activated, ramps in the Str are disrupted. Ramping activity in the Str was found to scale across multiple time intervals when the hold period duration is changed. For example, when the period between the start of a trial and the action cue is changed (increased from 3 seconds to 12 seconds), ramp duration would also increase to 12 seconds. This suggests that these ramps result from the hold period rather than other behavioural artifacts.

Mirzaei et al. (2017) analysed data from correct Go trials of several combined datasets where rats performed the SST task (Leventhal et al., 2012; Schmidt et al., 2013; Mallet et al., 2016). In the GPe, ramps were present in around 71% of neurons before a Go cue, with 47% positive and 52% negative. In the population, there was a net average positive ramping. However, this was likely due to the amplitude of the positive ramps being around four times higher. Similar ramping patterns were found for STN, with ramps present in around 77% of units, with 44% positive and 55% negative. The STN is densely and reciprocally connected to the GPe. STN typically excites the GPe, and also receives inhibition from it. Therefore it is possible that ramps in the STN may lead to GPe ramps. Mirzaei et al. (2017) suggested that the ramps may relate to a braking signal, which stops premature movements (where it is initiated too early), as previously suggested by Frank (2006).

In all of the above studies, ramp detection relies on firing rates average over multiple trials. It is currently unclear how well the identified ramps in the trial-averaged data capture ramping at a single trial level. In the context of cortical accumulation of evidence, it has been debated whether ramping trial-averaged responses in the motor cortex are ramping or stepping in their single trials (Latimer et al., 2015; Zhao and Kording, 2018; Zoltowski et al., 2019). This is important because it directly affects possible functions of ramping activity (Narayanan, 2016).

The slope and the maximum amplitude of the ramp can both be used to encode timing in the brain. Steps do not have slopes (or only very sharp slopes) such that the amplitude does not vary over time. For example, the amplitude can have many varying values during a ramp, but during a step it could be either the firing rate before or after the step only. This slope and varying amplitude values can allow for more complex encoding such as timing in the brain. In dopamine models, ramping has also been proposed to be related to the reward prediction error (Lerner et al., 2021). This would not be possible with only stepping in single trials. Furthermore, it is important to ascertain whether the trial-averaged ramp is ramping or stepping in single-trials, as many steps in single trials can add up to a ramp. If the ramp is made up of many steps it may point to differences in neural information processing compared to ramping single-trial activity.

In terms of ramp detection, supervised models require labelled data to fit the models. As ramps are continuous, their characteristics must be considered in order to label them. Supervised methods are common in other fields, such as wind energy production. For this, many datasets have been manually labelled and thus methods have been developed to tune parameters to labelled datasets (Gallego-Castillo et al., 2015). In contrast, unsupervised methods are used in neuroscience as we have no clear definition or function of a ramp. For example, in the wind ramp literature, the characteristics of the ramp can be determined by how they affect energy production. Thus, ramps that affect energy production can be labelled. However, in neuroscience, it is not clear how ramping activity affects behaviour. Thus it is difficult to identify relevant characteristics of ramps in order to label them for supervised methods.

Latimer et al. (2015); Zoltowski et al. (2019); Zhao and Kording (2018) used unsupervised methods to distinguish between ramping and stepping in single trial data. While Zoltowski et al. (2019) found that neurons in many datasets were

found to be ramping non-linearly or stepping, Zhao and Kording (2018) used cross-validation and found, using a similar dataset as Zoltowski et al. (2019), that neurons prefer ramping models instead.

Previous studies have found ramping behaviours in different basal ganglia regions in various tasks (Mirzaei et al., 2017; Donnelly et al., 2015; Emmons et al., 2017; London et al., 2018). However, few studies have looked at single trials in the basal ganglia, and multiple regions at the same time (some investigated two regions but not four or five). We need an algorithm that is able to distinguish between ramping and stepping single trials to be able to test whether population-averaged ramps are built out of steps or ramps. This would affect the mechanism by which inhibition is represented. Understanding inhibitory behaviour can help explain movement disorders and possibly pave the road for future treatments. Including more realistic inhibitory behaviour mechanisms could also improve current biologically-inspired agents. Furthermore, the developed ramp identification algorithm could also be used to characterise ramps more widely in the brain and help understand temporal expectations and information in other regions. Finally, we ask whether ramping behaviour can be observed in various basal ganglia regions and if these trial-averaged ramps are also reflected as ramps in their single-trial activity. Can ramping activity explain action restraint and behavioural inhibition in the brain?

2.4 Ramp detection

In the above studies, quantitative approaches to ramp detection include fitting a linear model (see Emmons et al. (2017); London et al. (2018)) or principal component analysis (Mirzaei et al., 2017; Donnelly et al., 2015). Four main characteristics usually define ramps; namely, the amplitude (power), the length (duration), the slope, the direction and the time of occurrence (Gallego-Castillo et al., 2015), see fig.2.4. The direction can either be a positive ramp when the firing rate increases over time, or a negative ramp when the firing rate decreases over time. Ramps are increases (or decreases) in firing rate over a significant period of time. If there is no minimum duration, then any increase in value can be considered a ramp.

It is difficult to properly assess the effects of ramping firing rates without first being able to characterize ramping activity (length, amplitude, slope, and direction).

One field where ramp detection algorithms are useful and have been extensively




Figure 2.4: Ramping properties. Adapted from (Narayanan, 2016).

developed and tested is in the generation of power via wind turbines (Gallego-Castillo et al., 2015). A consistent change in wind strength can be defined as a ramp, and the ability to accurately characterize, predict and forecast them enables efficient energy capture and prevents wind overload. However, wind ramps in different locations throughout the world vary significantly, making it difficult to construct a consistent definition of what constitutes a ramp (Gallego-Castillo et al., 2015). For example, Hannesdóttir and Kelly (2019) found differences in ramp characteristics amongst sites in southern Sweden and Denmark.

For wind power, most research has utilised a binary system to classify the ramps (Gallego-Castillo et al., 2015). This means that every point of the time series is assessed and classified as either ramping or not ramping. The key elements considered most important for ramp definitions are the thresholds for amplitude (power) and time. A segment of the time series is considered as a ramp when the criteria for these thresholds are met (e.g. amplitude and duration of the ramp are above a predetermined value). However, this binary definition has been challenged by Gallego-Castillo et al. (2015), who instead suggested that a more continuous metric, that captures the characteristics of a ramp should be implemented. They used wavelets to detect ramps across various characteristics. Wavelet methods were also developed in (Hannesdóttir and Kelly, 2019; Cheneka et al., 2020). Nevertheless, one issue with wavelet techniques is that it tends to select for features we are not

interested in. For most wavelets, ramps that reverse in direction after hitting the peak instead of plateauing are favoured, due to the nature of wavelet shapes.

There are several considerations in designing a good ramp detection algorithm. First, the number of hyperparameters should be minimized without affecting the robustness of the results. Second, the speed of execution, which affects computation time, is also a key consideration when dealing with large datasets. For example, in wind ramping data, there are often many time points sampled from periods of months to years. Similarly, with neural spiking data, sampling is down to millisecond periods, with many hours of data typically collected over several sessions. Lastly, the accuracy of the algorithm is also important. Several considerations should be made here, such as whether it correctly identifies ramps in the specified environment and if it appropriately filters and smooths out any noise.

Evaluating whether a ramp detection algorithm successfully meets the above criteria is not straightforward. Typically, two key kinds of evaluation metrics are utilised. First, the error between simulated ramps and whether the algorithm accurately detects these can be quantified using metrics such as Root-mean-square-error (RMSE). Second, tests of precision and accuracy for whether ramps are accurately captured or not. However, one difficulty of ramp detection algorithms is comparing one penalty or regularization hyperparameter to another.

Here we introduce several main algorithms used to detect ramps, including l_1 -trend filtering and swinging door algorithms from the wind ramp literature. We also introduce changepoint detection methods developed in mathematical research fields that can be adapted to ramp detection, for example, the pruned exact linear time (PELT) and opt algorithms.

2.4.1 l_1 -sliding window (l_1sw)

One of the methods developed to help solve the problem of detecting and characterising ramp events during wind power generation is the l_1 -sliding window algorithm (l_1SW) (Sevlian and Rajagopal, 2013). The method was used to build statistical models of wind ramps so that electric grids can be optimised to capture as much wind energy as possible. Their method required three primary steps to detect ramps in a large time series. Here we briefly describe the stages, with further detail of each given in the next section. First, l_1 -trend filtering (l_1tf) gener-

ates a piecewise solution of the large time series. Trend filtering is typically used to remove ‘noise’ in the time series by smoothing it out. Second, each ‘piece’ or segment in this solution is then evaluated and subject to a binary classification of either being a ramp or not being a ramp. Briefly, this evaluation looks at the minimum amplitude, rate of the ramp, slope, and the difference between the minimum and maximum observed amplitude within the segment (see fig. 2.4). Each segment was then given a ramp score, and segments that failed the above thresholds were given a score of 0. These segments will then be merged and provided with a score. They then used dynamic programming to find ramps with the highest scores. Lastly, the computation time is significantly reduced by using a sliding window. A sliding window separates the l_1 tf-filtered signal into multiple overlapping intervals and processes them parallel. The amount of overlap is determined by the maximum length of the ramp of interest. Overall, l_1 SW is effective at detecting and characterising ramps and is commonly used as a benchmark for new methods (for example, in Cui et al. (2021)). The loss function for l_1 tf is defined as (eq. 2.1):

$$\hat{x} = \min_{\hat{x}} \frac{1}{2} \|x - \hat{x}\|_2 + \lambda \|D\hat{x}\|_1 \quad (2.1)$$

where x is the original time series, \hat{x} is the predicted piece-wise linear time series, D is the second derivative operator and $\lambda \|D\hat{x}\|_1$ helps to implement sparsity in \hat{x} . λ is the regularization hyperparameter, where higher λ values indicate greater regularization. Regularizing hyperparameters ensures that models minimize the adjusted loss function while avoiding overfitting or underfitting. $\|x - \hat{x}\|_2$ denotes the l_2 norm of the difference between the original time series and the predicted piece-wise linear time series. $\|D\hat{x}\|_1$ is the l_1 norm of the second derivative of the predicted time series.

When λ is at its maximum, the piecewise component generated from the time series will be a single line. When λ is at its minimum ($\lambda_{\min} = 0$), the piecewise component converges to the original data. Different values of λ can result in different lengths of selected ramps. The maximum value of λ (eq. 2.2) is obtained by multiplying the second derivative operator D by the transpose of the second derivative order D^T and multiplying the inverse of that result with the

first derivative of the original time series Dx .

$$\lambda_{\max} = \left\| (DD^T)^{-1}Dx \right\|_{\infty} \quad (2.2)$$

The sliding window is a commonly used algorithm to reduce processing time via parallel processing. The piecewise signal generated by `l1tf` above is segregated into multiple overlapping windows. The number of windows M is decided by the window length (W_L) and window overlap (W_O). The index of each window is (i). The index of the start of each window, w_{si} is given by eq. 2.3 and the end index of each window, w_{ei} is given by eq. 2.4:

$$w_{si} = (i - 1)(W_L - W_O) \quad (2.3)$$

$$w_{ei} = (i)W_L - (i - 1)W_O \quad (2.4)$$

where subsets of \hat{x} are converted into the form $X_{w_{s1}:w_{e1}}, \dots, X_{w_{sM}:w_{eM}}$. X denotes the whole time series. The ramp detection method then determines whether each segment in a window is a ramp or not by applying rules such as minimum amplitude, power and slope. The results are then combined, and the longest ramp is selected in the sections where windows overlap.

2.4.2 Swinging door algorithm (SDA)

At about the same time that `L1SW` was developed, the SDA, which was initially designed as a compression algorithm, was adapted to be used in wind ramp detection (Florita et al., 2013). SDA utilises a single hyperparameter, ϵ , to define the width of the 'door' that swings outwards from its pivots, as shown in fig. 2.5. A higher value of ϵ makes the algorithm less sensitive to noise, similar to how higher values of λ in `l1tf` increase time series regularisation. A lower value of ϵ makes the algorithm more sensitive to small fluctuations in the time series. Using fig. 2.5 as an illustration, we run through an example interaction of the SDA:

(1) The hyperparameter ϵ and the first data point together determine the pivot of the two hinges. Using this, the SDA will draw a line from each of the hinges to point A (see grey lines in fig. 2.5A). This is the first "swing" of the door. The

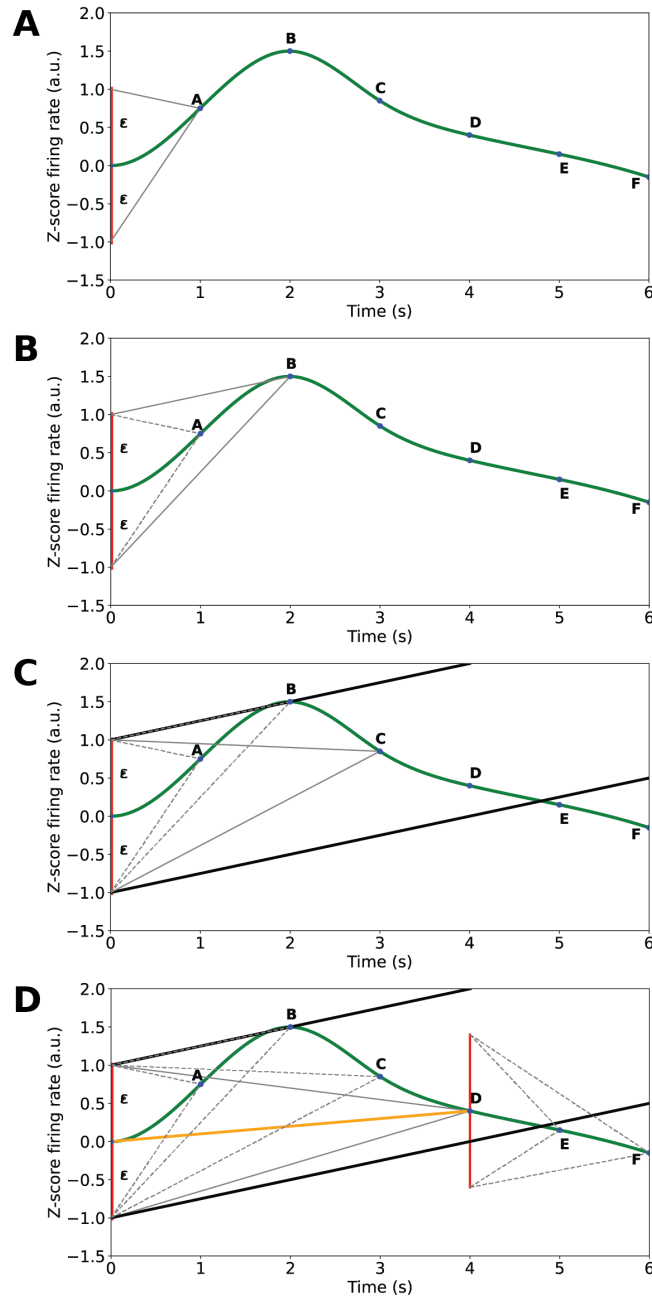


Figure 2.5: The SDA for the extraction of ramps from the time series, the scale is arbitrary for explanation purposes. Green curve shows interpolation with several time points highlighted (A-F). Red line extends on the y-axis in both positive and negative directions from zero to the value of hyperparameter ϵ , here, 1. Orange line in D shows the piecewise fitted component following one cycle of the algorithm. Adapted from Florita et al. (2013).

algorithm then goes sequentially through the time series.

(2) On the next swing, the doors are then pivoted to point B (grey lines to B in fig. 2.5B). As the 'triangle' for point A is within the 'triangle' for point B, there is no inflection point, and the doors will all be newly drawn to point B.

(3) On the next swing to point C (fig. 2.5C), the 'triangle' for point B is not within point C's 'triangle', thus reflecting an inflection point where only the lower door is able to extend to point C. Both doors have to be parallel to each other at their widest.

(4) On the final swing to point D (fig. 2.5D), the lines are almost parallel and are unable to cross to any further point without swinging the door too wide. Thus, a piecewise component is then formed from the first point to D (orange line in fig. 2.5D). This starts a new cycle of the algorithm, where point D will be the new starting data point for the next iteration.

To formalize, let us take A_c to be the amplitude of the current point while A_l and A_u to be the amplitudes of the lower and upper bounds of the swinging door. A_g and A_s are the amplitudes of the gate points and startpoint of each SDA segment. t_c , t_g and t_s represent the current time, gate time, and start time, respectively. As long as P_c satisfies eq. 2.5:

$$A_l < A_c < A_u \quad (2.5)$$

$$A_u = A_s + \frac{(A_g + \Delta E - A_s)(t_c - t_s)}{t_g - t_s} \quad (2.6)$$

$$A_l = A_s + \frac{(A_g - (\Delta E - A_s))(t_c - t_s)}{t_g - t_s} \quad (2.7)$$

the point is within the hinges of the swinging door, and the next point will be checked again. The advantage of the SDA is that there is only a single parameter: e . Larger values on e might result in significant ramp events being filtered, while a smaller e may result in capturing small fluctuations.

2.4.3 Changepoint algorithms

L1SW and SDA are both ramp detection methods. Here, we introduce changepoint detection algorithms which attempt to detect points in a time series where statistical properties change (Killick et al., 2012). Unlike l1SW, where the piecewise components were built before determining the changepoints, we could use changepoints algorithms to first determine the changepoints before building the piecewise components. Let us define our time series as $x_{1:n} = x_1, \dots, x_n$ with a number of changepoints, m , where the timepoint of each changepoint is $p_{1:m} = p_1, \dots, p_m$. The position of the changepoints are integers between 1 and $n - 1$ where $p_1 > 0$ and $p_m < n$. We can also define $p_0 = 0$ and $p_{m+1} = n$. Each i th segment with this set of changepoints will be split into $x_{p_{i-1}+1} : p_i$. Changepoint detection aims to estimate the indices in $p_{1:m}$, where the number of changepoints, m , may or may not be known depending on the problem we are solving. Here, the total number of segments would be $m + 1$. To identify changepoints, we often attempt to minimize (Killick et al., 2012):

$$\min \sum_{i=1}^{m+1} [C(x_{(p_{i-1}+1):p_i})] + \beta f(m) \quad (2.8)$$

where $C(\cdot)$ is a cost function for each segment and $\beta f(m)$, is the penalty or regularization factor to help prevent overfitting, where $f(m)$ is the penalty function and β is the penalty constant. Many changepoint algorithms can be framed as having a cost function, a search method, and a constraint on the number of changepoints (Truong et al., 2020). However, this categorization does not include Bayesian methods such as the Hidden Markov Model. The cost function of a changepoint algorithm informs us how homogeneous the time series is. A signal that is similar throughout with few changepoints would have a low cost, whereas a heterogeneous signal with many changepoints would have a high cost. The search method of a changepoint algorithm is used to obtain an exact or approximate solution to eq. 2.8. The value of the penalty would affect the number of changepoints detected. Selecting a high penalty would result in few changepoints, and selecting a low penalty would result in a high number of changepoints.

2.4.4 Opt

Opt is a search method based on dynamic programming to efficiently compute the solution to detect a fixed number of changepoints (Jackson et al., 2005) (see Algorithm 1). Opt uses a linear penalty which means that

$$\beta f(m) = \beta m \quad (2.9)$$

Using eq. 2.9 above, we obtain:

$$\min \sum_{i=1}^{m+1} [C(x_{(p_{i-1}+1):p_i}) + \beta] \quad (2.10)$$

For a time series $x_{1:T}$, we define the set of all possible vectors of changepoints to be $P_T = \{p : 0 = p_0 < p_1 \dots p_m < p_{m+1} = T\}$ and we define $F(T)$ to be 2.10 as follows:

$$\begin{aligned} F(T) &= \min_{p \in P_T} \left\{ \sum_{i=1}^{m+1} [C(x_{(p_{i-1}+1):p_i}) + \beta] \right\}, \\ &= \min_t \left\{ \min_{p \in P_T} \sum_{i=1}^m [C(x_{(p_{i-1}+1):p_i}) + \beta] + C(x_{(t+1):n}) + \beta \right\}, \\ &= \min_t \{F(t) + C(x_{(t+1):n}) + \beta\}. \end{aligned}$$

This allows us to recursively solve for $T = 1, 2, \dots, n$, giving us the minimal cost for the time series $x_{1:T}$. The cost of solving this for time T is linear in T .

Algorithm 1 Algorithm Opt.

Input: signal $\{x_t\}_{t=1}^T$, cost function $C(\cdot)$, penalty value β

Initialise: $n =$ Length of time series, set $F(0) = -\beta$ and changepoint output, $CPO(0) = \text{NULL}$. **For:** $p^* = 1, \dots, n$

Calculate $F(p^*) = \min_{0 \leq p < p^*} [F(p) + C(x_{(p+1):p^*}) + \beta]$

Let $p^1 = \arg\{\min_{0 \leq p < p^*} [F(p) + C(x_{(p+1):p^*}) + \beta]\}$

set $CPO(p^*) = (CPO(p^1), p^1)$

Output: set of changepoints in $CPO(n)$

Adapted from Killick et al., 2012

2.4.5 Pruned Exact Linear Time (PELT)

PELT is a changepoint search method that is exact with a linear computational cost (Killick et al., 2012) (see Algorithm 2). The exactness means that it finds the optimal solution to a problem, rather than an approximate, or near optimal, one. When we are unable to constrain the number of changepoints and the objective function to minimize the penalized sum of costs, PELT is a fast method to search for linear penalties. PELT extends the original Opt algorithm, whereby a pruning step is added within the dynamic program of Opt (Jackson et al., 2005). This reduces computational cost without affecting the algorithm's exactness. The pruning primarily helps remove values of cp which cannot be minima at each iteration of algorithm Opt, while still finding the global minimum of the cost function. This pruning rule is given in eq. 2.11 where there exists a constant K such that for all $t < s < T$,

$$C(x_{(t+1):T}) \geq C(x_{(t+1):s}) + C(x_{(s+1):T}) + K \quad (2.11)$$

Then, if eq. 2.11 holds, at a future time $T > s$, then t cannot be the last changepoint prior to T :

$$F(s) \leq F(t) + C(x_{(t+1):s}) + K \quad (2.12)$$

Algorithm 2 Algorithm PELT.

Input: signal $\{x_t\}_{t=1}^T$, cost function $C(\cdot)$, penalty value β , constant K .

Initialise: $n = \text{length of time series}$, set $F(0) = -\beta$, changepoints $CPO(0) = \text{NULL}$, $R_1 = \{0\}$ **For:** $p^* = 1, \dots, n$

Calculate $F(p^*) = \min_{p \in R_{p^*}} [F(p) + C(x_{(p+1):p^*}) + \beta]$

Let $p^1 = \arg\{\min_{p \in R_{p^*}} [F(p) + C(x_{(p+1):p^*}) + \beta]\}$

set $CPO(p^*) = (CPO(p^1), p^1)$

set $R_{p^*+1} = \{p^* \cap \{p \in R_{p^*} : F(p) + C(x_{(p+1):p^*}) + K < F(p^*)\}\}$

Output: Record changepoints in $CPO(n)$

Adapted from (Killick et al., 2012)

2.5 Spike-phase entrainment

In the above sections, we identified that ramping firing rates occur throughout the brain, and these ramps may have functions related to behaviour, such as action restraint. These inhibitory behaviours have also been found in oscillatory field potentials (Knyazev, 2007).

Local field potentials (LFPs) are fluctuations of voltage in extracellular neural activity and can be recorded via an extracellular microelectrode. LFPs are usually studied using time-frequency based analysis using information such as phase, amplitude and frequency. They can be found in many regions in the brain on a large spectrum of scale measured through electrophysiology in small circuits or wider, larger EEG signals (Wilson et al., 2018). LFPs reflect pathological activity in the brain, and differences in activity can be found in many neurological diseases such as Parkinson's, schizophrenia and autism (Uhlhaas and Singer, 2010; Engel and Fries, 2010; Boutros et al., 2015). Furthermore, the study of these LFPs has also shown to improve treatments such as deep brain stimulation (Priori et al., 2013; Rosin et al., 2011). It has also been suggested that oscillatory activity of LFPs allows brain regions to communicate with each other and transfer information across regions (Fries, 2005). Behaviourally, the phase of oscillatory activity before stimulus onset can be used to predict responses (Spaak et al., 2014; Busch et al., 2009) and beta oscillations in the motor system related to the maintenance of the status quo (Engel and Fries, 2010). Overall, LFPs can be used to study neurological diseases, behavioural states, and behavioural responses.

Still, some sceptics believe oscillations are epiphenomenon or have no mechanistic relevance to the function of the brain. There are several reasons that this is unlikely. First, studies show that extracellular signals could influence firing in the cortex (Anastassiou et al., 2011), and more recently, in the cerebellum (Han et al., 2018), this is known as ephaptic coupling. More importantly, our understanding of the brain is far from complete; thereby, claiming that LFPs are irrelevant is far-fetched. Some have argued that oscillations are simply a direct product of spikes. However, they can also be caused by other non-action potential induced currents, such as calcium waves that arise as a nonlinear function of spikes and firing rates. Lastly, oscillations still provide a useful readout of neural activity and could help improve our understanding of behavioural states.

Extracellular potentials, such as LFPs, originate from the sum of transmembrane currents across neurons and therefore do not have a single source. Distant sources could be volume conducted from other regions as electromagnetic fields can propagate throughout the brain. In addition, the structures of dendrites and axons and the timing of synaptic inputs can determine which frequency bands are reflected in the electrode. In general, these LFPs reflect the organization of large-scale brain dynamics (Pesaran et al., 2018), but more specifically, they stem from a variety of electrophysiological processes such as synaptic inputs and spiking. However, more often than not, neurons do not generally use field potentials to communicate, and field potentials are effects rather than causes of neuronal transmission (Buzsáki et al., 2012).

Usually, researchers are interested in the low-frequency component of the field potential ($<100\text{Hz}$), which is generally believed to be dominated by synaptic activity. However, it is difficult to interpret as anatomy, neuronal connectivity, nonsynaptic ionic currents and tissue conductance can influence it (Herreras, 2016; Buzsáki and Wang, 2012; Lindén et al., 2011, 2010; Cole and Voytek, 2018). Furthermore, spiking activity could still contaminate low-frequency components up to 10Hz (Waldert et al., 2013). While these low-frequency components are unlikely to be contaminated ($<10\text{Hz}$), and high-frequency components are more likely to be contaminated, the extent to which this affects analysis is not entirely clear (Waldert et al., 2013). Ultimately, we should proceed with caution when analysing LFPs as there could be many sources of influence. The extent to which these different sources affect the field potential is still largely unknown, and progress in understanding this would rely largely on building better forward models (Pesaran et al., 2018). When constructing models of LFP generation, one can generate model-based benchmarking data to test how electrode placement, volume conduction and other synaptic components will affect the field potential. This will ultimately lead to better electrode array designs, sensors, and clinic treatments for disorders (Pesaran et al., 2018).

The relationship between LFPs and spiking activity has been studied extensively. The same electrode is often used to record both the appropriate LFP and a single-unit recording. If spikes preferentially occur during one or more certain phases of the oscillation, the firing of a neuron is said to be entrained to an LFP oscillation. This relationship is typically defined by both the phase and frequency of the LFP

where spikes occur, and the consistency of that occurrence. One common method to study this relationship includes examining the spike-triggered average (STA) of the LFP and attempting to predict the occurrence of spikes from LFPs and histograms of spike-LFP phase-locking (Zanos et al., 2011). Spike-field coherency (SFC) is another measure used to determine the phase synchronisation between the spike trains and LFPs by providing information on the preferred phase where spikes occur on the LFP.

There are several contending hypotheses for the function of SFC in brain circuits. Fries (2015) suggests that SFC in the gamma band plays a role in the communication between brain regions and can inform us about the strength of inter-region connectivity. Others, however, have suggested that it better reflects excitation-inhibition interactions as gamma frequencies tend to be very sensitive to low-level stimuli (Ray and Maunsell, 2015; Ray, 2022; Bartoli et al., 2020). Some studies have shown that SFC can provide information about the encoding of the environment. More specifically, SFC has had various behavioural roles in different regions. For example, in the hippocampus, phase information from the theta range can help provide information on the animal's position in the environment (Jensen and Lisman, 2000). In the striatum, fast-spiking interneurons could be crucial for selecting between different actions and processing information (Berke, 2009). In the cortex, Womelsdorf et al. (2014) found that burst spikes were locked to beta and gamma frequencies across the anterior cingulate and lateral prefrontal cortex.

Differences in cell types in entrainment occur frequently. For example, nearly all neurons in the hippocampus preferentially fire at some phase of the theta oscillation. However, there are cell type disparities in the phase of the gamma oscillation entrainment (Klausberger and Somogyi, 2008). In addition, cells preferentially match their firing in the orbitofrontal cortex with either the theta or gamma oscillation, but not both (van Wingerden et al., 2010). Some of these cell-specific variations in the frequency-selectivity of entrainment most likely result from connection variations, while some are more clearly brought on by inherent frequency preferences (Wilson et al., 2018).

Spike removal from LFPs is usually done by lowpass filtering the time series with a cutoff of 300-500Hz. This is important to ensure that the spikes do not influence the shape of the waveform. Although many still assume that this is enough to remove spiking activity, spikes tend to bleed through to the LFP despite lowpass

filtering (this is generally termed the spike-bleed-through effect). This may affect LFPs up to frequencies of 10Hz, and it is theoretically possible they affect even lower frequencies (Waldert et al., 2013). When studies do not attempt to remove spikes from LFPs, results can be misleading and result in invalid and flawed conclusions.

Zanos et al. (2011) found artifactual correlations between LFPs and spikes via ground truth simulations. These correlations can distort common spike-phase coupling results, as the coupling is due to the spike instead of the actual coupling of oscillations and spikes. Thus to prevent these artifacts after low-pass filtering, they developed a spike removal algorithm based on Bayesian principles. Using this algorithm, they were able to remove spurious frequencies above 100Hz from the spike-triggered average wave. Furthermore, this method also prevents SFC problems. Waldert et al. (2013) proposed a method to detect spike contamination in LFPs, and showed that the characteristics of the spike, such as its duration, amplitude and firing rate, can influence the extent of the contamination on the LFP. Despite the numerous studies on the spike-LFP artifacts, none have looked closely at how waveform differences can affect the phase of spike-phase coupling.

It is common to use spikes and LFPs from different electrodes to prevent spike bleed-through in the LFP. The argument is that extracellular action potentials cannot be detected by electrodes separated by more than $400\ \mu\text{m}$. Any further would be unable to pick up spiking activity while still able to capture the gist of lower frequency data. Ray (2015) highlighted two key issues of this. Firstly, spikes between neurons are often correlated up to 3mm away with high temporal precision (Smith and Kohn, 2008). This distance is far further than the $300\ \mu\text{m}$ of separation between electrodes that studies typically use. Thus, if an action potential occurs on one electrode, there is a high likelihood of an action potential event on the nearby electrodes. Second, the oscillations can vary significantly across electrodes which are separated just $400\ \mu\text{m}$ apart, thus rendering the SFC spurious and inaccurate. For example, in primary visual cortex (V1), gamma rhythms at $400\ \mu\text{m}$ had significantly different frequencies (Ray, 2015).

Chapter 3

Quantification of firing rate ramps

3.1 Chapter overview

In this chapter, we investigate algorithms to help detect firing rate ramps accurately and robustly, namely `l1tf`, SDA, Opt and PELT. The first two algorithms, `l1tf` and SDA, are taken from the wind ramp literature. Note that the `l1tf` we used in this chapter is not the same as `l1SW` (Sevlian and Rajagopal, 2013). We did not include a score function in our algorithm, but only used `l1tf` to generate a piecewise solution. Opt and PELT are changepoint algorithms which have been used to detect the transition between tasks in functional magnetic resonance imaging (fMRI) (Saggar et al., 2018). We developed a method of detecting ramps by first implementing a changepoint search algorithm, Pruned Exact Linear Time (PELT), and then constructing a piecewise solution from the changepoints to determine whether a segment is classified as a ramp.

One key problem with ramp detection algorithms is the sensitivity of results to hyperparameters. More specifically, in `l1-trend filtering` (`l1tf`), the λ parameter and the γ parameter both influence the ramp detection, and we show that small changes in these parameters can result in large changes in the detection of the ramps. Similarly, in SDA, the parameter ϵ can also affect results. Lastly, PELT also has a penalty hyperparameter, P . We investigated the effects of the different algorithms' hyperparameters on ramp detection and quantified the robustness of each method with respect to their hyperparameters.

An accurate ramp detection algorithm has to consider the statistics of the data they are analyzing. Therefore, we first generated realistic simulations of firing rates and then compared the different methods. We then identified the methods that best detect ramps in our firing rate data, which can be characterized as a non-stationary univariate time series. Although there are multiple papers investigating firing rate ramps in neuroscience, both in the context of single trials and mean firing rates of all trials, there has been no investigation into the effectiveness of any ramp detection methods.

We then developed a quality measure specifically for ramps in firing rate data that favoured ramps with a long duration and a high signal-to-noise ratio. This quality is key to determining whether ramps exist in our empirical dataset in the following chapter 4. Finally, we test our quality measure with different firing rate scenarios such as uniform, step and ramp firing rates to ensure that it is well-suited to accurately identify them in neural activity.

We show that PELT has the highest precision in detecting the ramps compared to the other methods. We also show that our quality measure helps us select firing rate ramps with a good signal-to-noise ratio and are less likely to be steps in firing rate, where the increase in activity ends quickly.

3.2 Methods

3.2.1 Algorithms

In this section, we will discuss the four main algorithms from the wind-ramp and changepoint literature, namely, the SDA, `l1tf`, PELT and Opt. While the first two algorithms are used to detect ramps with an unknown number of changepoints, Opt helps to find the exact solution where there are a known number of changepoints. We have included this to compare our algorithms to the 'ideal' scenario where we have more constraints on the algorithm. As discussed in the background chapter 2, these algorithms approach the ramp detection problem in different ways. `l1tf` generates an approximate piecewise solution given a penalty, and we only extract ramp statistics such as changepoints, length and amplitudes from the piecewise solution. PELT and Opt, on the other hand, are changepoint detection algorithms. Thus we begin by first finding the changepoints and then

'drawing' the ramp from the changepoints. Finally, SDA attempts to build a piecewise solution from only a single hyperparameter timepoint by timepoint. Whilst the equations and pseudocode for these algorithms were given in chapter 2, here we provide the values for the hyperparameters of each.

Hyperparameters for simulations.

We define each piecewise component from one point to another as a segment. Each of these segments are filtered by a minimum amplitude of 0.5 z-score, and a minimum length of .35s. All l1tf parameters used are shown in table. 3.1

Method	Parameter	Value
l1tf	λ	500
l1tf	γ	.0001
SDA	ϵ	.1
PELT	Penalty	25
OPT	no. of changepoints	2
All	Length _{min}	.5 seconds
All	Amplitude _{min}	.5 z-score

Table 3.1: Parameters for simulated ramps.

3.2.2 Firing rate simulations

The statistical structure of neuroscience data is different from wind ramps; thus, having realistic firing rate simulations to test the reliability of each method is useful. To test the various algorithms, we generated piecewise ramps where we have the 'ground truth' changepoint, endpoint and amplitude of the ramp. In our simulations, we used the parameter values shown in table 3.2. The ramp length was chosen as this was the average hold period in the data sets we applied the ramp detection algorithms to in the following chapter 4. Similarly, both the baseline firing rate and slope are realistic representations of our data. After generating the piecewise ramp, we add noise drawn from a Gaussian distribution with standard deviation levels between 5 and 50. All scenarios were run over a period of 4 seconds, simulating two seconds before and after an aligned event, such as a go cue or a stop-signal. Our simulated data is then converted to z-scores. This procedure is similar to how we analyzed our experimental data later in chapter 4.

Parameter	Value
Base firing rate	5 spikes/second
Slope of ramp	25 spikes/second
Length of ramp	1 second
Length of time series	4 seconds

Table 3.2: Parameters for simulated ramps.

One of the key features of a good ramp detection algorithm is to be able to clearly detect differences between non-ramps, such as step-like functions and constant functions, compared to ramping ones. Thus we generated four main firing rate simulations to test our algorithms: the uniform firing rate, a firing rate with a step in it, a ramp and a piecewise ramp with Gaussian noise applied (See fig. 3.1). In the first three scenarios, we generated 500 poisson spike trains that were modulated by different firing rates (uniform, step and ramp). This roughly corresponds to the number of single trials we have from one session from each neuron in the data sets we use in chapter 4. For the last scenario, we generated a piecewise ramp and added Gaussian noise on top; this is to add a 'ground truth' comparison with less complicated statistics of the data. Similar to the firing rate simulations above, all scenarios were run over a period of four seconds, simulating two seconds before and after an aligned event, such as a go cue or a stop-signal. This is similar to how we analyzed our experimental data later in Chapter 4. The average firing rates of the uniform and step processes were approximately 33 Hz, and for the ramp processes, around 40 Hz. The firing rates in fig. 3.1 demonstrate one example simulation for each scenario.

3.2.3 Quality

We developed a quality measure of how strong the ramps are. This was done for several reasons. First, some detected ramps have less variation in firing rates, while others have more. Ramps with less variation (with respect to the detected ramp) are preferred, as these have a higher signal-to-noise ratio. Second, we want to prevent spurious ramps that are noisy. Third, we would like to penalize shorter ramps to distinguish them from steps. Finally, the quality measure also helps us to better differentiate between steps in the firing rate and ramps in the firing rate.

We define the startpoint of a putative ramp to be a and the endpoint to be b , such

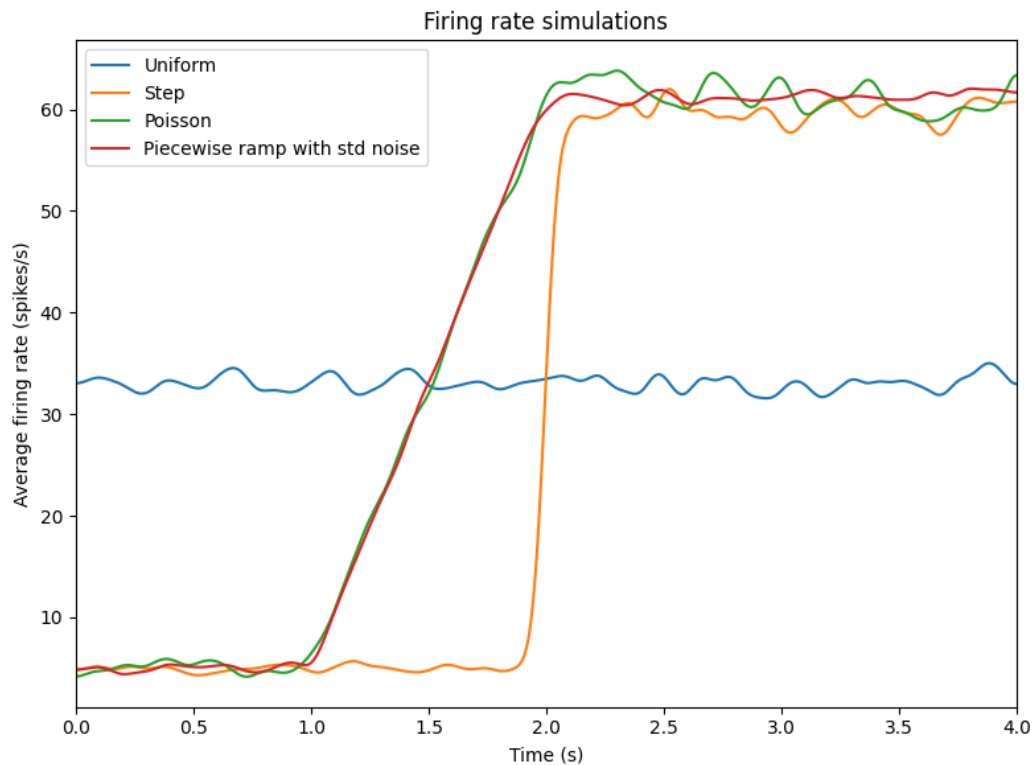


Figure 3.1: Examples of simulated firing rates in four scenarios. All simulations, except the piecewise ramp, are the average of 500 poisson spikes train modulated by a particular firing pattern. In the uniform simulation, there is a constant firing rate of 31 spikes/s. The step simulation begins with an average firing rate of 5 spikes/s and sharply increases to 60 spikes/s at 2.0(s). The poisson ramp simulation begins with an average firing rate of 5 spikes/s and starts ramping to 60 spikes/s at 1.0(s). The piecewise ramp simulation with Gaussian noise has similar parameters to the poisson ramp but is modulated by Gaussian noise with a standard deviation of 5.

that $x_{a:b}$ is the z-score firing rate during the time of the putative ramp. We define the piecewise segment generated out of the startpoints a and b to be $y_{a:b}$. $\sigma_{x_{a:b}-y_{a:b}}$ is the standard deviation of the difference between the original signal, $x_{a:b}$ and the piecewise ramp, $y_{a:b}$. $x(b)$ is the z-score amplitude at the endpoint, and $x(a)$ is the z-score amplitude at the startpoint.

The signal to noise ratio (SNR) is defined as:

$$SNR = \frac{(x(b) - x(a))^2}{\sigma_{x_{a:b}-y_{a:b}}^2} \quad (3.1)$$

$$f(x) = \begin{cases} x & \text{if } 0 < x < 1 \\ 1 & \text{if } x > 1 \end{cases}$$

$$Q = SNR * f(b - a) \quad (3.2)$$

We tested the quality on the four different types of simulated data: uniform, step, poisson ramp and a ramp with Gaussian noise. This was to ensure that the quality measure was able to distinguish between stepping and ramping.

3.2.4 Evaluation metrics

We require measures of performance to evaluate the effectiveness of the various ramp detection algorithms on firing rate ramps. We have adapted metrics from the changepoint detection field, which are also typically used for evaluating supervised learning algorithms (Aminikhanghahi and Cook, 2017; Gallego-Castillo et al., 2015). We begin by defining some key terms in the confusion matrix in table. 3.3.

	Classified as ramp	Classified as non-ramp
True ramp	True positive (TP)	False negative (FN)
False ramp	False positive (FP)	True negative (TN)

Table 3.3: Confusion matrix for ramp detection

The ramp classification we employ above does not include the accuracy of the detection. For example, if two ramps were detected, but only one ramp exists, the

TP would be 1, and the FP would be 1, and if one ramp was detected and only one ramp exists, the TP would be 1, and the FP would be 0. Lastly, if no ramps exist, the TP and FP would both be 0. False negatives occur when we detect fewer ramps than the ground truth, and false positives occur when we detect more ramps than the ground truth.

We use the precision metric to evaluate the occurrence of confusion matrix outcomes. Precision is defined as:

$$\text{precision} = \frac{TP}{TP + FP} \quad (3.3)$$

It is useful to identify how many detected ramps are classified correctly as ramps. In addition, accuracy helps to identify how often both our ramps and non-ramps are correctly classified:

$$\text{accuracy} = \frac{TP + TN}{TP + FP + FN + TN} \quad (3.4)$$

Finally, we used root mean square error (RMSE) to determine the error of our detected ramp. The RMSE is the square root of the sum of the squared errors for both the changepoint and the endpoint of the ramp divided by two:

$$\text{RMSE} = \sqrt{\frac{(CP_{\text{groundtruth}} - CP_{\text{detected}})^2 + (EP_{\text{groundtruth}} - EP_{\text{detected}})^2}{2}} \quad (3.5)$$

Here, we use CP for a ramp's startpoint (or changepoint) and EP for the endpoint of a ramp. We later show the average RMSE of all true selected ramps. This means that only in situations where the number of ramps is accurately detected that we include the RMSE into the calculation.

3.3 Results

3.3.1 Sensitivity of γ parameter in `l1tf`

To evaluate each method quantitatively, we have used measures such as RMSE and precision in the later part of the results. In this section, we show some examples of our individual simulations and each method to better illustrate the advantages and disadvantages of each method. These examples are selected to highlight problematic ramp detection, more standard examples can be found in Appendix A.

To obtain the changepoints for `l1tf`, we have to set a threshold for γ values, as `l1tf` provides an approximate piecewise solution. However, there is no clear threshold that obtains perfect changepoints for every piecewise solution. As a result, `l1tf` tends to provide more changepoints than the other methods, some very close to each other (see the cluster of changepoints around .9 sec in fig. 3.2). Also, in this particular example, we can see two changepoints between 1.8s and 2.0s; this makes the endpoint selection of the ramp inaccurate as the second changepoint should definitely be included as a ramp. This is often a result of `l1tf` being able to detect very small changes in slope to produce the ideal piecewise solution. PELT, however, produces changepoints, which helps to solve the problem of sensitivity to small changes in slope.

To better quantify the threshold γ used to obtain changepoints, we calculated the maximum threshold value we could use to capture a ramp in 95% of simulations (see red vertical line in fig. 3.3). Using this threshold is strict and would mean that 5% of simulations will not capture a ramp (this is worst than the currently selected threshold of γ in this noise level and will be shown later in the chapter in fig. 3.6). Furthermore, using this threshold means that only five simulations will have two changepoints, while most will have between three to nine changepoints. Although γ is not comparable to the penalty in PELT (a more similar comparison would be λ), we ran a similar analysis looking at the number of changepoints across PELT while varying the penalty. For a similar noise level, PELT algorithm correctly provided two changepoints for $> 99\%$ of the simulations.

The original SDA algorithm tends to suffer from ‘bumps’ in the time series. These bumps reside outside of the swinging door for very few timesteps but cause a

new piece to be generated from the algorithm. This problem is again evident as the SDA example generated a changepoint in the middle of the ramp (see fig. 3.2). Both PELT and l_{1tf} do not have this problem where 'bumps' significantly affect their ramp detection.

3.3.2 PELT penalty hyperparameter is the most robust

Most ramp detection algorithms will require a hyperparameter that determines the extent the solution is regularized. An ideal scenario would be to have as few hyperparameters as possible, while producing accurate results that are robust to small changes in the hyperparameter. Here, we test the sensitivity of the results to tuning the hyperparameter in l_{1tf} and PELT. For l_{1tf} , we chose a range of λ values from 300-700, which equates to about a 40% upper and lower bound from the baseline. We chose a range of penalty values from 15-35 for PELT, with the original value being 25. This equates to about a 40% upper and lower bound from the baseline.

Figure 3.4 shows an example where l_{1tf} is unstable across different values of λ . A short ramp is detected at low regularization levels, but the ramp then disappears and reappears around a λ value of 420.

To better quantify the extent of this problem, we calculated the RMSE between each level of regularization for both l_{1tf} and PELT. If the detected changepoints for a ramp between two regularization levels are exactly similar, the RMSE will be 0. For example, in fig. 3.4, the PELT detection for this particular simulated ramp will yield only zero values for the RMSE. In this case, it will result in 19 zeroes as we compare between 20 adjacent values of the penalty. The l_{1tf} , however, will have several non-zero values between the λ levels that vary in ramp changepoint and endpoint. The distribution of RMSE between different regularization levels across 100 simulations shows that PELT has a lower count of high RMSEs than l_{1tf} (see fig. 3.5).

Figure 3.5 shows the distribution of RMSE between different levels across 100 simulations. We have removed the results where there is a less than .01s RMSE as these mostly constitute zeroes or very minor differences in changepoints. The l_{1tf} is much less robust to changes in the hyperparameter than PELT. The overall count of RMSE above .1 is higher than PELT.

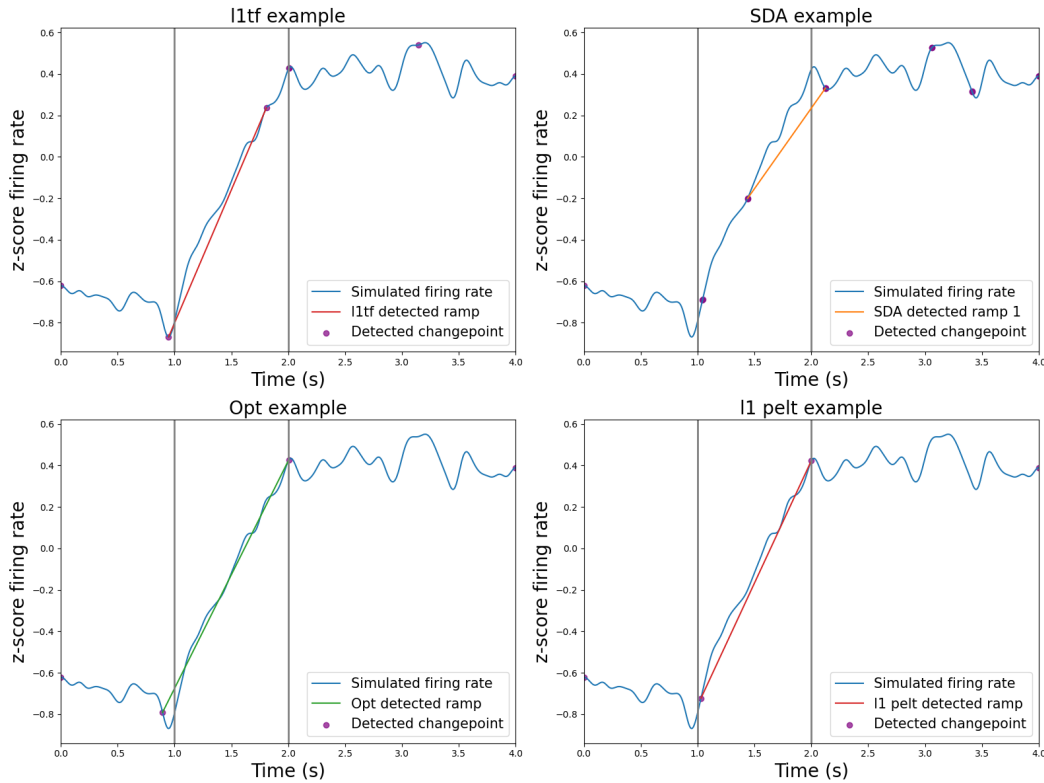


Figure 3.2: More changepoints were detected in l1tf and SDA. Each panel shows the detected ramps and detected changepoints of each method. The changepoints of each method are shown by a purple circle marker. We used a simple piecewise ramp with Gaussian noise (standard deviation of 35ms) where the ground truth changepoint occurs at 1 second, and the endpoint is at 2 seconds. In the l1tf example, a λ value of 500, and a γ value of 0.001 was used. In the SDA example, ϵ was set to .1. With PELT, a linear penalty of 25 was used. In the Opt example, there is no hyperparameter, but we have specified the number of breakpoints, here two. A linear cost function was also used. This should act as an ideal ‘ground truth’ algorithm. Upper left panel shows the detected ramps using l1tf. Upper right panel shows the SDA algorithm applied to the same simulation. Lower left panel shows the Opt changepoint algorithm with the number of breakpoints given. Lower right panel shows the PELT algorithm.

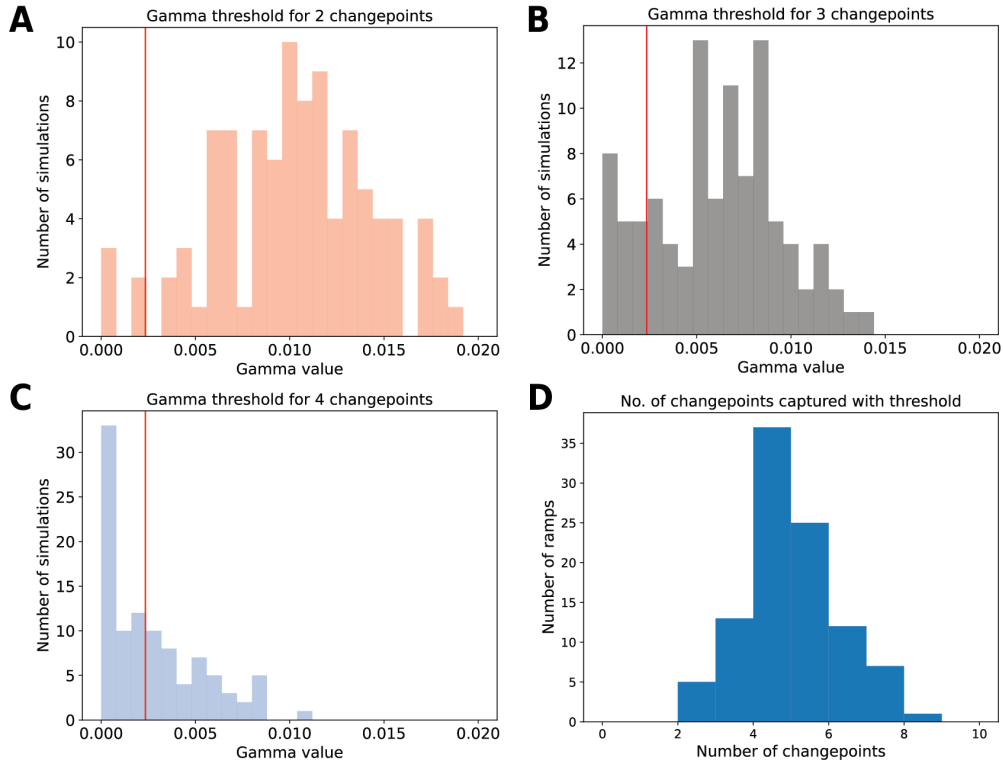


Figure 3.3: γ threshold will often capture a wide number of changepoints. A. Distribution of γ thresholds for the algorithm to output two changepoints. These simulations were run with 30ms standard deviation Gaussian noise similar to fig. 3.2. B,C. Similar histogram but for three and four changepoints, respectively. Red vertical line in panels represent threshold required to ensure that 95% of simulations will capture at least one ramp. D. Number of changepoints detected when the threshold is applied. For most simulations the threshold still yields multiple changepoints.

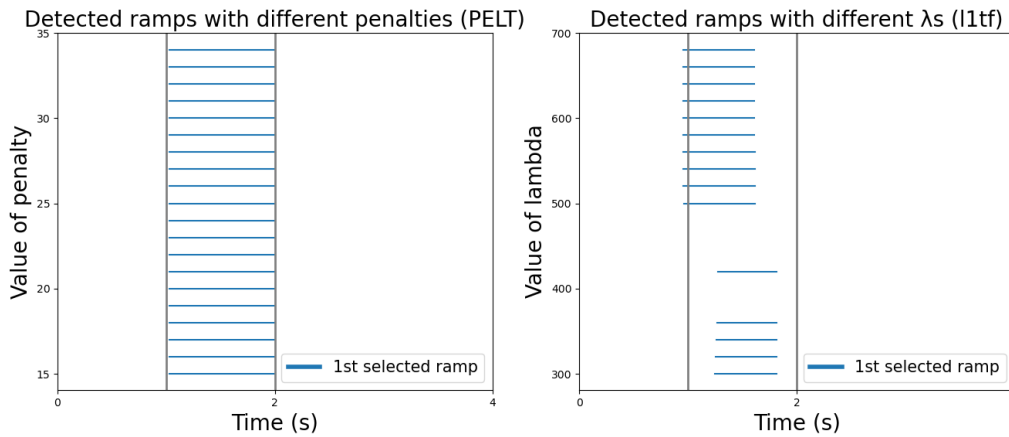


Figure 3.4: lambda is more inconsistent in selecting ramps. Left panel shows the selected ramps at different levels of the hyperparameter linear penalty for PELT. Right panel shows the selected ramps for the different values of lambda. Grey lines represent the 'ground truth' changepoints.

3.3.3 PELT outperforms other methods in precision of ramp detection.

To test how Gaussian noise affects the precision of the ramp detection, we measured the precision of each method while varying the standard deviation of the Gaussian noise applied to a piecewise linear ramp. PELT performs best on precision at most noise levels, and its performance is comparable to the 'ground truth' of the Opt algorithm (see fig. 3.6). The l1tf performs well at lower noise levels but begins to 'miss' ramps at higher noise levels. SDA has poor performance throughout most noise levels compared to the other methods. At very high noise levels, the precision of PELT method begins to get better than even that of the Opt method. One possible explanation for this is that the penalty factor helps to smooth the data and deals with noise better, but the Opt solution will overfit on noisy data. The average precision is not only higher for the PELT method, but the precision itself also has a lower standard deviation, meaning that there are fewer simulations where PELT is not precise.

3.3.4 PELT shows higher RMSE than l1tf on detected ramps.

To measure how accurate each method is in detecting changepoints across varying levels of noise, we measured the RMSE of the already detected ramp. Only ramps that were detected were included in this analysis; thus, this analysis does not rep-

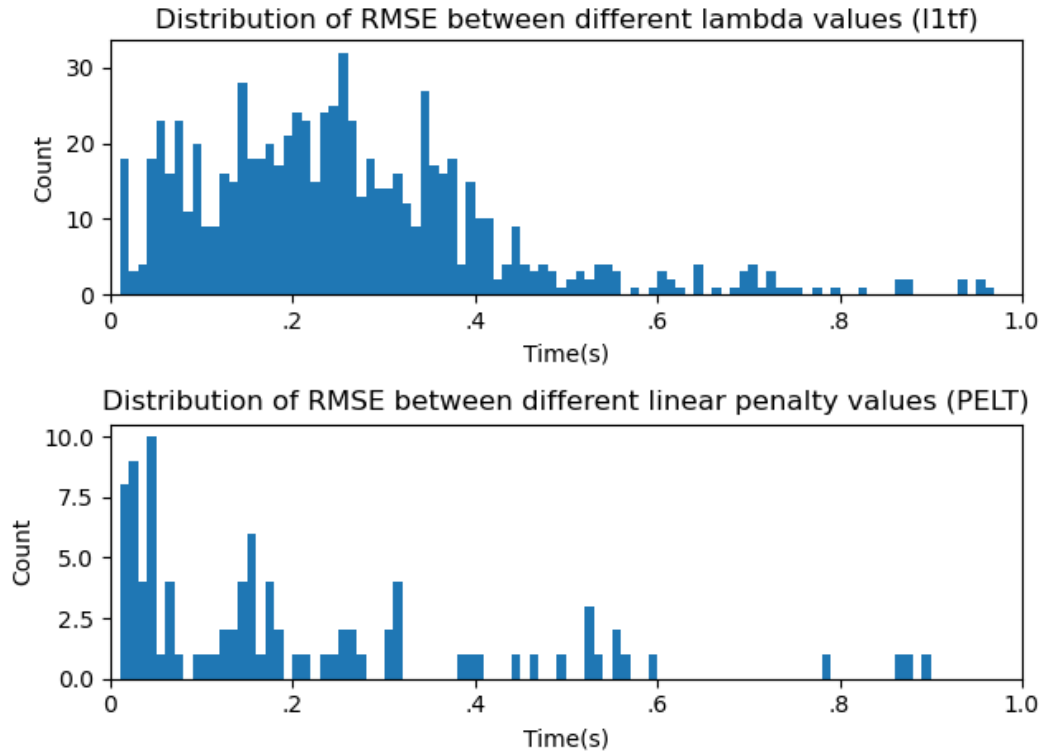


Figure 3.5: The l1tf is much less robust to changes in the hyperparameter compared to PELT. Top panel shows the distribution of RMSE errors between all levels of λ . Bottom panel shows the distribution of RMSE errors between all levels of penalties in PELT. We have removed the results where there is a less than .01s RMSE as these mostly constitute zeroes or very minor differences in changepoints. We can see that the difference in RMSE between lambda values is much higher than the difference between linear penalties of PELT, making PELT a more robust method.

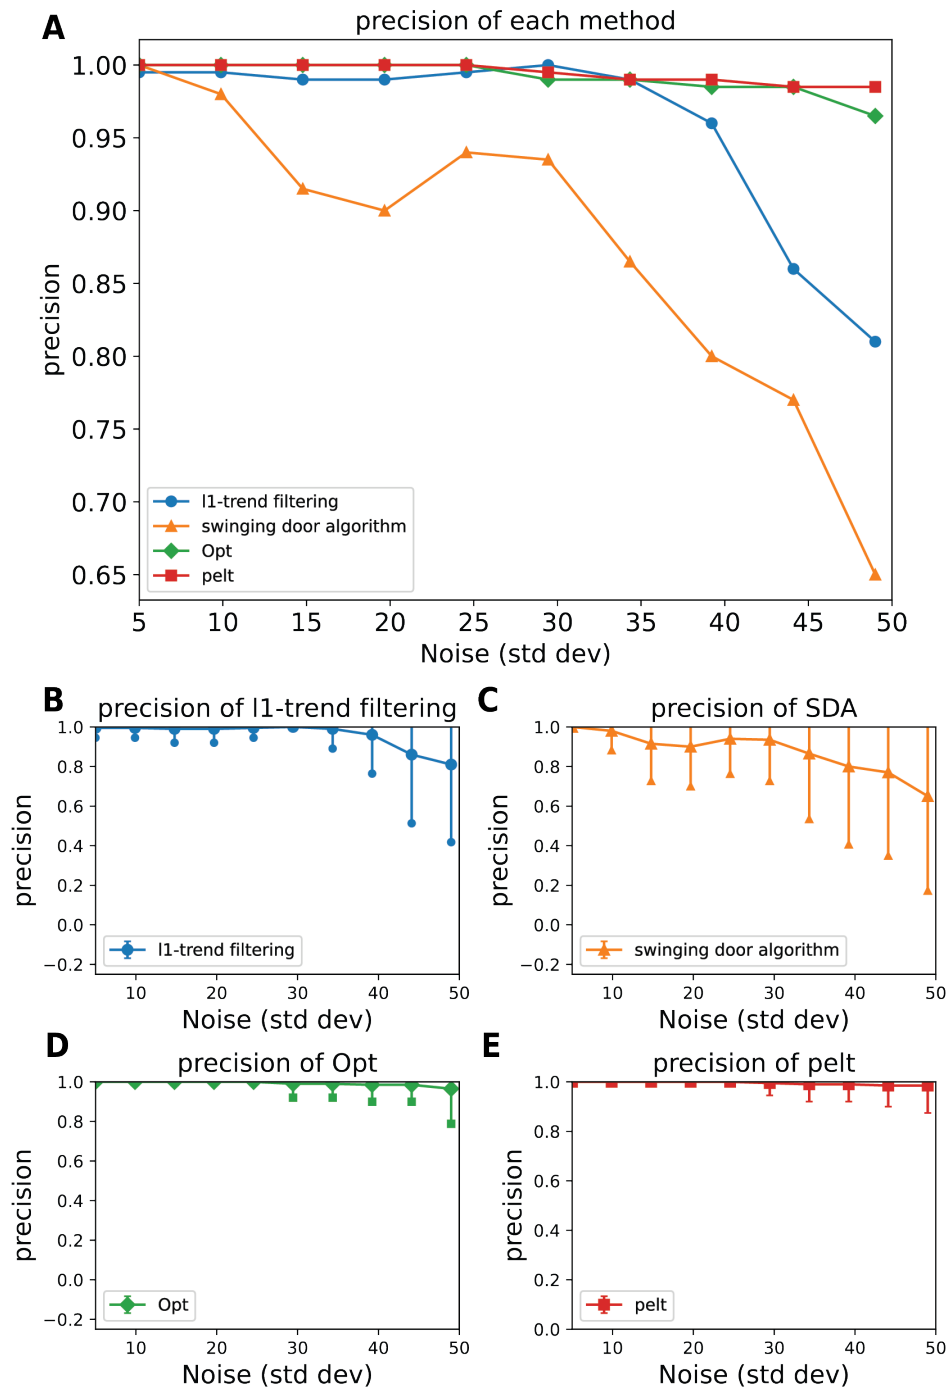


Figure 3.6: PELT has the highest precision compared to other methods. A. Precision of each method across varying noise levels (5-50ms standard deviation). B-E. Precision of each method across varying levels of noise with standard deviation of their precision plotted. PELT has the lowest standard deviation of precision.

resent missed ramps. SDA performs the worse throughout all levels (see fig. 3.7). PELT performs better than l1tf at lower noise levels but worse at higher noise levels. This means that, at low noise levels, PELT has both higher precision and better accuracy (RMSE) than l1tf. At higher noise levels, PELT can still detect the ramps (higher precision) but may be less accurate in detecting the changepoints due to the noise. We can see from the standard deviations of RMSE in fig. 3.7 that l1tf rarely 'accepts' ramps above a RMSE of .2s.

3.3.5 l1tf is biased towards shorter ramps.

To further investigate the strengths and weaknesses of each method, we look at the distribution of detected ramp characteristics (changepoint, endpoint, amplitude, length) of each method.

The distribution of lengths, changepoints and amplitudes were evaluated for the simulation with 30 standard deviation Gaussian noise. l1tf seems to detect the changepoint relatively accurately compared to other methods. However, it detects a shorter length in general. This suggests that l1tf tends to fail to detect the endpoint accurately.

3.3.6 Quality measure in simulations

We verified our method by looking at the different qualities of each simulated scenario. We wanted to ensure that the quality measure found the difference between steps, uniform firing rate and a ramp. We chose to generate poisson firing rates from a perfect ramp. We then generated the other simulations (step, uniform, and piecewise) with low noise level of 5 to replicate the noise in the generated poisson firing rate ramp. Furthermore, we wanted to capture ramps where the signal-to-noise ratio was higher, where the detected piecewise component better represents the actual firing rate. For the non-uniform scenarios, there is usually only one low quality segment (see fig. 3.9). The step firing rate scenario has three segments, two segments where the firing rate is uniform and one segment where the firing rate is a step. In the CDF, we could see that the step simulated data has a couple of segments below 25 quality and some between 25-50, which is the step (see fig. 3.10).

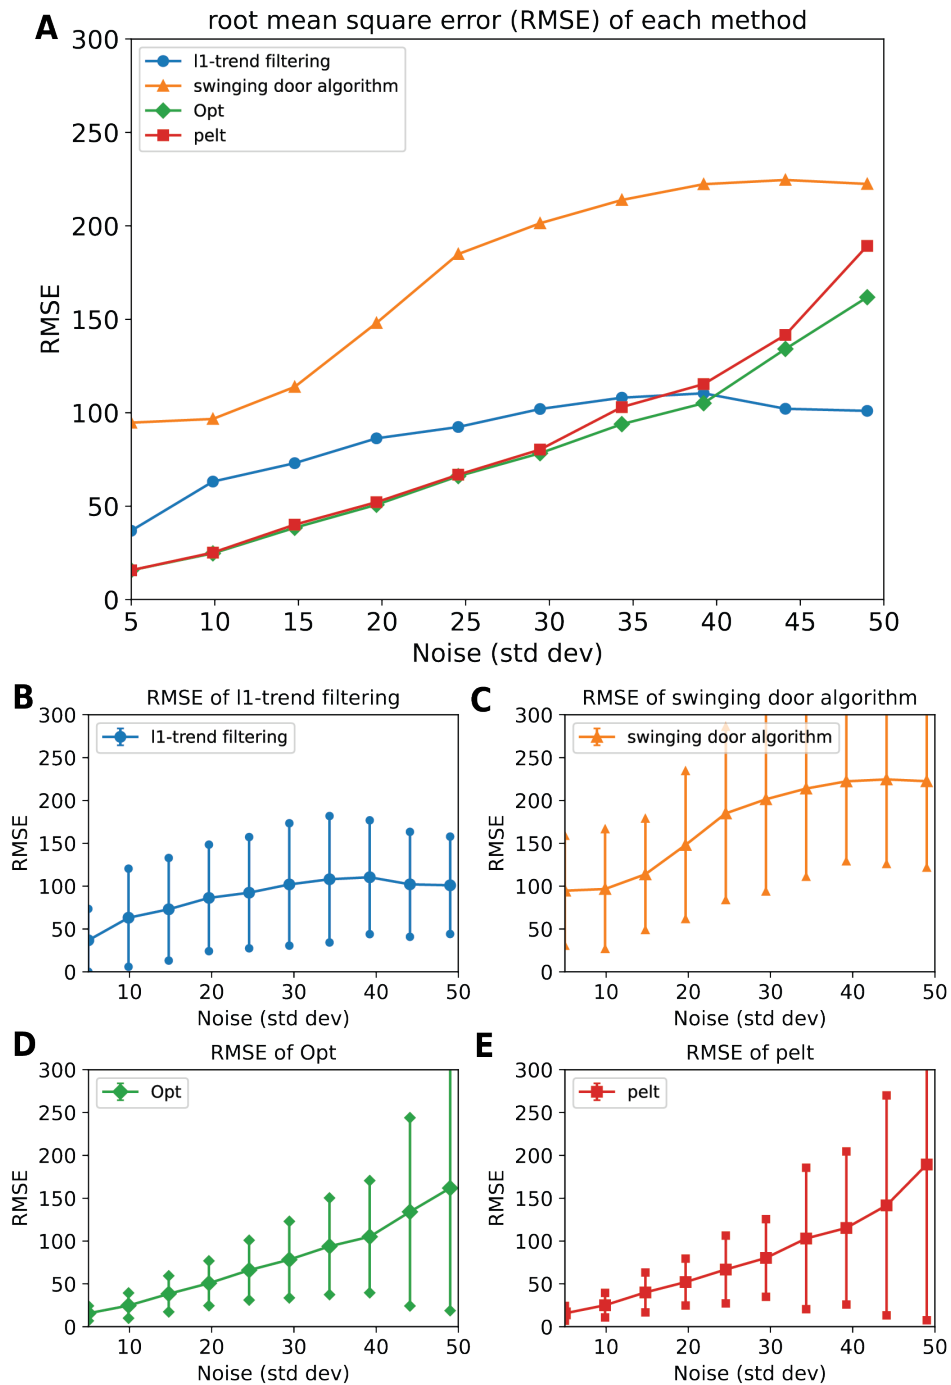


Figure 3.7: RMSE for each method. A. Root mean square error of each method across different noise levels. B-E. RMSE of each method across varying noise levels with a standard deviation of their RMSE plotted. Each simulation was run 1000 times. At low noise levels, PELT has the lowest RMSE, although most methods have relatively low RMSE at low noise levels. At high levels, PELT and l_1 -trend filtering performs the worst in terms of RMSE, but that is simply because it can detect more ramps.

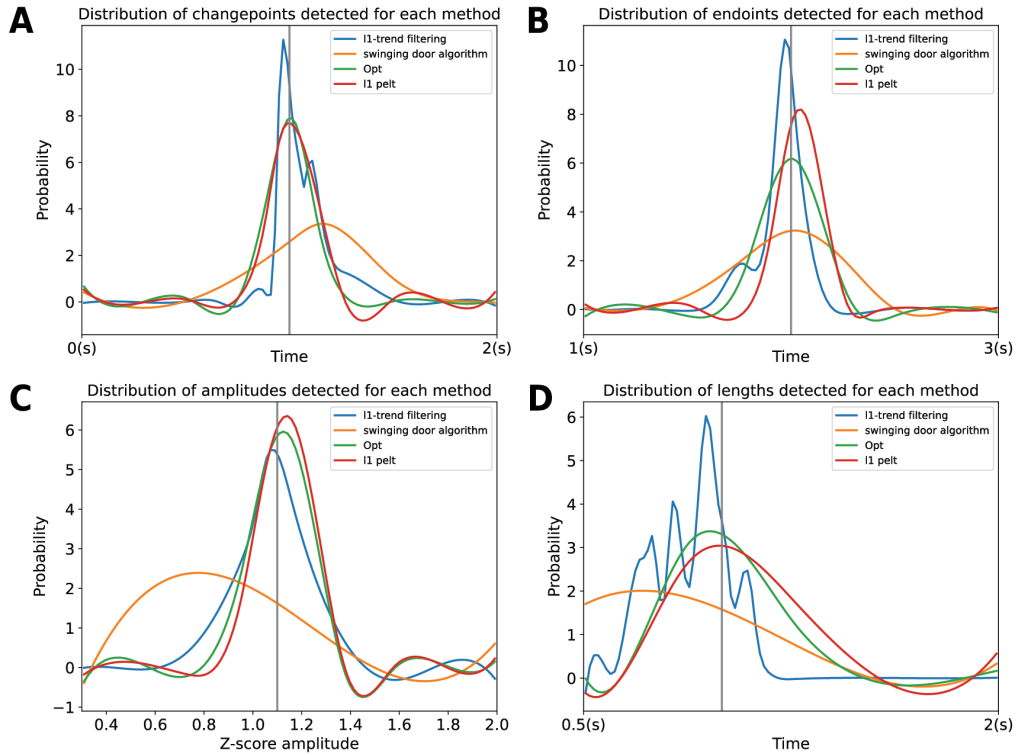


Figure 3.8: l1tf detection is inconsistent. Results shown are run with 30ms standard deviation Gaussian noise similar to fig. 3.2. A. Probability distribution of different changepoints for each method. Grey vertical line represents the ground truth changepoint. B. Probability distribution of different endpoints for each method. Grey vertical line represents the ground truth endpoints. C. Probability distribution of different detected amplitudes for each method. Grey vertical line represents the ground truth amplitudes. D. Probability distribution of different detected lengths for each method. Grey vertical line represents the ground truth lengths.

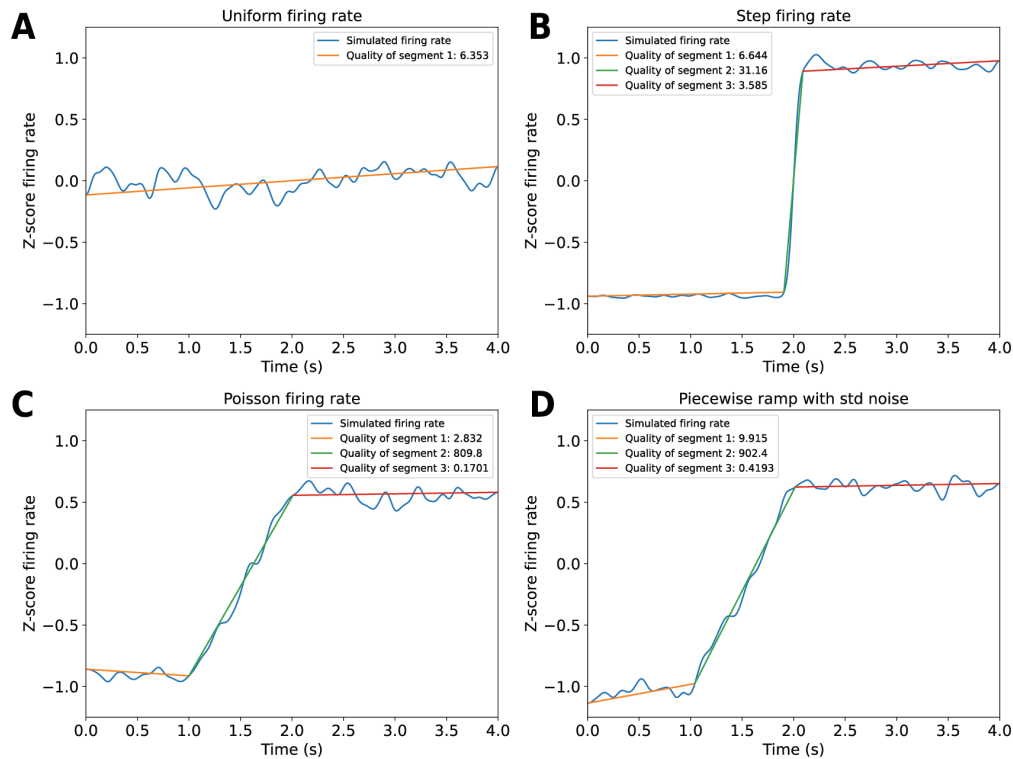


Figure 3.9: Quality between stepping segment and ramping segment is significantly different. Quality of each detected ramp in the caption. Each panel demonstrates the different types of z-score firing rate patterns. Gaussian noise of 5 was used to replicate noise level in poisson simulated firing rate. Each coloured line shows the detected piecewise segment using PELT. No amplitude or length filters were applied to the segments. Qualities of a step piece are low.

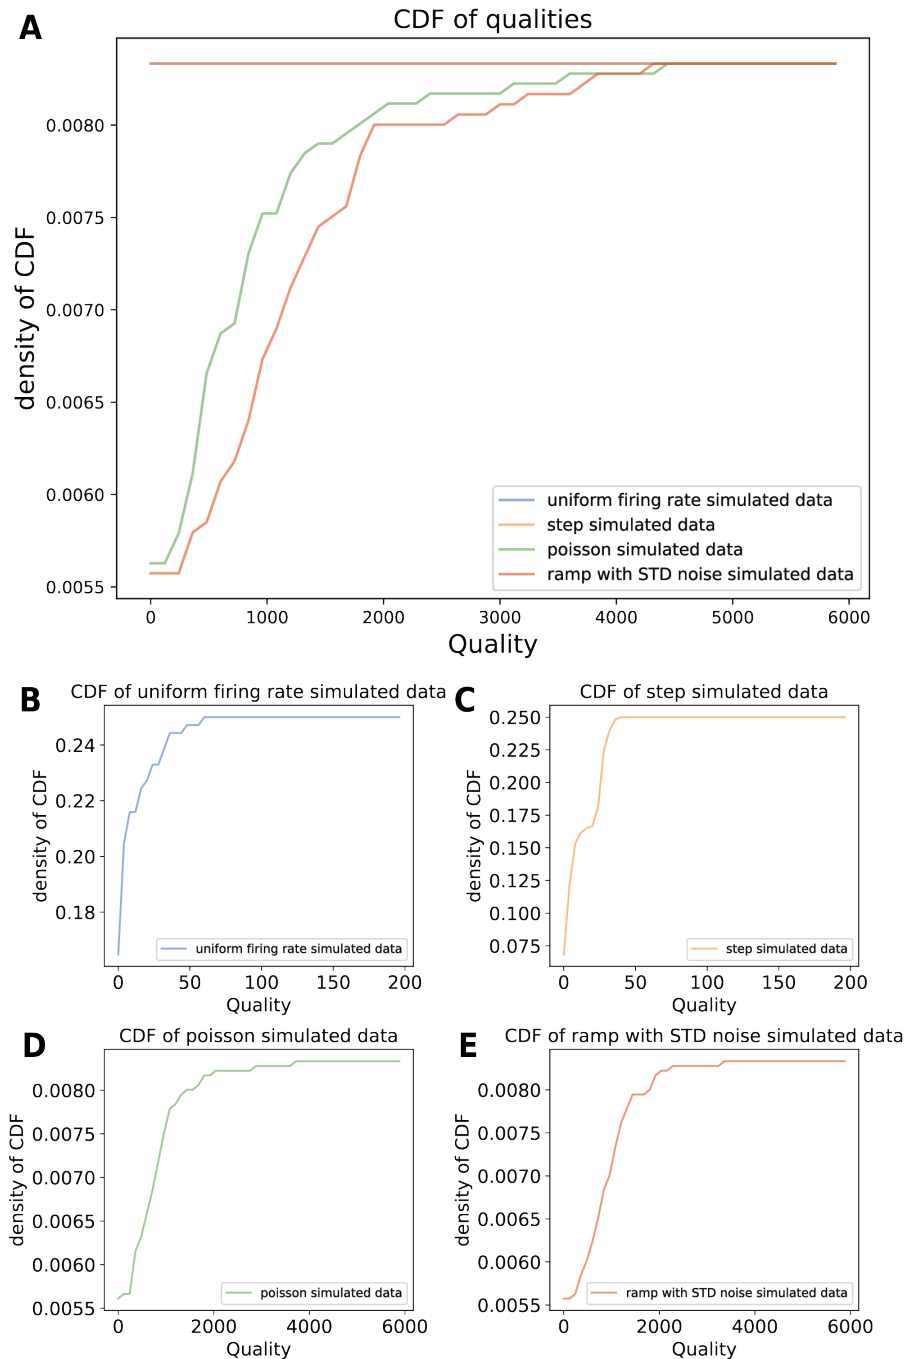


Figure 3.10: Qualities for step simulated data and uniform firing rate simulated data are often below 50 quality. The cumulative distributive function of quality on each of the different simulation scenarios. A. CDF of uniform simulated data. Qualities for noisy uniform firing rate are primarily under 50. B. CDF of quality for step simulated data. As with panel A, most of the qualities are also below 50. C. CDF of poisson simulated ramp. Most of the values are between 250 to 1500 for the ramp segments. D. CDF of piecewise ramp with Gaussian noise. Most values are above 250 with ramp segments.

3.4 Discussion

In this chapter, we investigated algorithms to detect and characterize firing rate ramps using realistic simulations of firing rates. We demonstrate that these algorithms have trade-offs with the number of hyperparameters and the algorithm's effectiveness. We also used PELT to build piecewise segments and then filter them for ramps. PELT is able to detect ramps more accurately, especially at higher noise levels. This method is also more robust to hyperparameters and requires one less hyperparameter than `l1tf`. Our method also uses a similar number of hyperparameters as SDA but is far more precise. Furthermore, our ramp detection also has a lower RMSE than SDA.

3.4.1 Comparison with previous approaches

Ramp detection algorithms should have as few hyperparameters as possible while still being able to detect ramps successfully. While SDA has only one easily tunable hyperparameter, its performance on ramp detection is generally poor. Furthermore, a ramp can be disrupted via a 'bump' that exceeds the door of the SDA algorithm for just one time point, making it very sensitive to 'bumps' in an otherwise perfect ramp. This makes it less suitable for noisy firing rate ramps.

Using `l1tf` to generate piecewise components performs better than SDA. However, `l1tf` requires two hyperparameters, λ and γ , which are sensitive and are not robust to small changes. λ at low values provides a better fit to the data but generates many more segments. In contrast, high values of λ generate fewer segments but provide a worse fit. We found that λ can affect ramp detection in unpredictable ways, making it difficult to select and not robust. Furthermore, γ values used to extract changepoints can be sensitive, and the number of changepoints can vary widely despite the data having similar properties (e.g. low Gaussian noise with a similar ramp).

In the `l1SW` (Sevlian and Rajagopal, 2013) and an optimized version of SDA (Cui et al., 2021), a ramp score function was used to help solve the problem of combining segments. In essence, each segment was given a score, and segments that did not pass the ramp thresholds were given a score of 0. Segments will then be combined and similarly given scores. The ramp with the highest score will then be selected. This method of combining segments helps to alleviate the problems of

both λ and γ parameters. However, combining segments in such a way increases the error between the signal and the detected ramp. This is not ideal in firing rate ramps where precise firing times of neurons can influence behaviour.

PELT uses a single robust parameter and performs the best amongst these methods across our simulations. We constructed segments out of the changepoints that PELT detected and filtered each of these segments by our ramp rules. In contrast, `l1tf` provides an exact piecewise solution given a regularization parameter λ , after which the inflection points of the solution have to be extracted as changepoints.

Wind ramps from the same region and within the same dataset tend to be very similar (e.g. wind ramps measured from a wind ramp farm in Norway have very similar structures (Hannesdóttir and Kelly, 2019)). However, in neuroscience, spiking activity within the same region from different neurons can be highly varied. For example, although we use z-scores to calculate our firing rates in later chapters, neurons within the same region may have different strengths of responses and fluctuations when their firing rates are ramping. This makes it particularly tough to tease out ramp events in firing rates.

Chapter 4

Properties of ramping activity in the basal ganglia

4.1 Introduction

4.1.1 Inhibitory control and ramping firing rates.

In the previous chapter, we investigated various ramp detection algorithms in firing rate simulations. We found that PELT performed the best for our simulations, and in this chapter, we applied it to firing rate data from basal ganglia. Although previous studies have identified ramping firing rates in the striatum or thalamus during hold periods (London et al., 2018; Emmons et al., 2017; Mirzaei et al., 2017; Donnelly et al., 2015; Narayanan, 2016), few studies have looked at ramping in multiple regions simultaneously. Here, we investigate the differences between ramping in several basal ganglia regions, including the STN, SNr, Str, GP and thalamus.

Previous studies used methods such as fitting a linear model (see Emmons et al. (2017); London et al. (2018)) or principal component analysis (Mirzaei et al., 2017; Donnelly et al., 2015). Neither of them provides the properties of ramps, such as the startpoint, endpoint, slope, duration, and amplitude. Here, we analyze these ramping characteristics in the 'hold' period of a stop-signal task (Schmidt et al., 2013; Mallet et al., 2016).

Typically, ramp detection relies on firing rates average over multiple trials. It is currently unclear how well the identified ramps in the trial-averaged data capture ramping at a single trial level. In the context of cortical accumulation of evidence, it has been debated whether ramping trial-averaged responses in motor cortex are ramping or stepping in their single trials (Latimer et al., 2015; Zhao and Kording, 2018; Zoltowski et al., 2019). This is important because it directly affects possible functions of ramping activity (Narayanan, 2016). Here we examined both trial-averaged and single-trial firing rates in rats during the hold period of a stop-signal task in order to determine whether ramping occurs in the context of action restraint and waiting.

Overall, we found ramping activity in neurons in all basal ganglia subregions and thalamus, and the properties of the ramps were mostly similar across subregions, except for the amplitude of the ramp. These strong ramps were found in the STN and SNr of the basal ganglia, supporting their role in action restraint. Furthermore, in neurons with ramping activity in the trial-average, single-trial activity exhibited either step or ramp-like firing rate changes. This indicates that ramps visible in the trial average while animals wait for a Go cue can be due to step-like or ramp-like single-trial activity.

4.2 Methods

4.2.1 Dataset

The dataset consisted of recordings from the STN, Str, SNr, GP, and the thalamus during a stop-signal task (Schmidt et al., 2013; Mallet et al., 2016). In the stop-signal task, each trial began with the animal entering an illuminated nose-port, where it waited then for an auditory Go cue. This time period constituted the hold period during which it had to remain still. The duration of the hold period was randomized between .8s and 1.2s. Data was taken from 10 rats, with an overall total of 400 sessions (each lasting 90 minutes). Only correct go trials were used.

4.2.2 Data Analysis

All data analysis was performed using custom-written MATLAB and Python scripts. The spiking data was aligned to either the Nose in or the Go cue event. Firing rates were then calculated by smoothing these values with a 45ms Gaussian kernel. We then obtained the z-score firing rate of each neuron using the session-wide mean and standard deviation of the firing rate of that neuron. For our ramp analysis, we looked at the firing rates from 2s before the Nose in event to 2s after the Go cue event.

4.2.3 PELT algorithm

We utilised the PELT changepoint algorithm, which was identified as the most accurate in the previous chapter 3 on simulated ramps. Following verification of the algorithm, we applied the method to detect ramps from firing rate data in recorded regions of the dataset. As we are mainly interested in ramps within the hold period, we only considered ramps that started within 0.5 seconds of the start of the hold period. When looking at the Go cue aligned trial-averaged firing rate ramps, we considered ramps within .5 seconds of where the possible Nose in event is. This is between 0.5 to 1.5 seconds before the Go cue.

Ramps with quality above 50 are less likely to be spurious ramps or steps as measured in chapter 3. Thus, we filtered away ramps with quality below 50. This is a conservative approach as the simulation conducted in chapter 3 used very low levels of Gaussian noise (see Appendix B for examples of firing rates where no ramps were found).

Single trial analysis

In the subset of neurons with a ramp detected in the average firing around the hold period, we analysed the single-trial spike trains to determine whether they reflected ramping or step-like changes in activity. This was done by calculating the likelihoods of a ramp and a step model fitting to the average firing rate of the detected ramp. For the step model, we built ten steps across the length of the ramp. These steps varied in length, determined by the slope of the firing rate. The time point of the initial step was determined to be at 10% of the total firing rate, and subsequent steps at every 10% increase in the amplitude of the firing rate (see

fig. 4.1).

We selected the timepoint of the step that has the highest likelihood 4.1 for likelihood of a spike train:

$$L = \prod_i \frac{p(V)^{x_a}}{x_a!} \exp^{-\Delta p(V)} \quad (4.1)$$

where x_a is the number of spikes in the a th bin, i is the total number of timepoints and V is the trial-average ramp.

We then attained the Akaike Information Criterion (AIC) score for the single trial for both the ramp and step model, defined as:

$$AIC = 2k - 2 \ln(L) \quad (4.2)$$

As the step model has a single parameter, the time step in 4.1, we defined $k = 1$. The ramp model has no parameters as it is completely constrained by the data. Thus we defined $k = 0$. L is the likelihood that the firing rate is derived from either the step or ramp model. The difference in AIC scores, δ , expresses how similar the experimental trial is to each of the models (see fig. 4.2). Sometimes, thresholds for δ values are used to determine strong or weak evidence between two models. However, in this case, they can be affected by the characteristics of the firing rate ramp. Thus, a better solution would be to compare these values with simulated values.

For each trial, we generated simulated Poisson spike trains from the average detected ramp of each neuron. In the step simulation, we provided the spike trains with ten steps (see fig. 4.1 for an example). In the ramp simulation, we simply generated ramps from the firing rate of the detected ramp. Finally, we matched the number of simulated trials with the number of actual trials in the dataset (see fig. 4.2). For each of these matched trials, a δ score is obtained. The result is a distribution of δ scores for simulated steps, stimulated ramps, and experimental data.

The distribution of δ scores can be affected by the characteristics of the detected ramp (e.g. length, amplitude, and slope) and model comparison methods are




Figure 4.1: Calculation of steps in the step model. The firing rate of the Step model (blue line) is fitted to an example detected ramp (black line). The time point of the first step was fitted to the point when the firing rate reached 10% of the maximum firing rate. Each subsequent step occurred when another 10% of the firing rate was reached so that a total of 10 steps correspond to the whole detected trial-averaged firing rate ramp. This method allowed us to account for firing rate changes in any ramp with steps, as it ensured that more steps occurred in regions when the firing rate increases were steeper.

frequently inconsistent when comparing models that differ greatly (Zhao and Kording, 2018); thus, we used these simulations to “baseline” δ score distributions. The full calculation process is shown in fig. 4.3. After the δ histograms were generated for the experimental and simulated data, the Kolmogorov–Smirnov (KS) test was then used to compare the distribution of δ scores within simulated model distributions and between them. For the stepping model, for example, KS scores were first calculated between all the generated δ histograms for stepping (step vs step). KS scores were then calculated between each of the δ histograms for stepping and ramping (step vs ramp). For each of these, the KS scores were averaged to produce an overall KS score for the step vs step and a score for ramp versus ramp. The difference between these two scores was then calculated (KS difference score). This process was repeated for the ramping data, with the final difference score between the averages of ramp vs ramp and ramp vs step. KS scores were also calculated for the experimental data by comparing the single experimental AIC histograms to all histograms for stepping and all histograms for ramping. A difference score was then calculated between the two, as with the simulated models. This process was repeated for all neurons in all regions, resulting in three scores for each neuron- the step, ramp and experimental KS difference scores. We will denote this KS difference score as Γ_i for each neuron i , such that positive values suggest the trials are more step-like and negative values suggest the trials are more ramp-like. Γ_i values close to 0 are unclear, and therefore not categorized as either ramping or stepping. Therefore, we categorized Γ_i values below -0.1 as ramping, and above 0.1 as stepping.

4.3 Results

4.3.1 Ramping activity occurs during the hold period throughout the basal ganglia.

In the stop signal task, a cue light indicated to the animal which nose port to enter. The entrance into the nose port triggered a hold period, during which the animal had to remain in the nose port until an auditory Go cue occurred. The frequency of the Go cue tone indicated whether the animal had to quickly move to the left or right side port (unless the stop-signal was presented as well). Notably, the animals had to remain still during the hold period to perform the task correctly

Algorithm 3 Algorithm ramping vs stepping.

Input: signal $\{x_t\}_{t=1}^T$, cost function $C(\cdot)$, penalty value β , constant K , minimum amplitude A , minimum length L .

Initialise: $n = \text{length of time series}$, set $F(0) = -\beta$, changepoints $CPO(0) = \text{NULL}$, $R_1 = \{0\}$ **For:** $p^* = 1, \dots, n$

Calculate $F(p^*) = \min_{p \in R_{p^*}} [F(p) + C(x_{(p+1):p^*}) + \beta]$

Let $p^1 = \arg\{\min_{p \in R_{p^*}} [F(p) + C(x_{(p+1):p^*}) + \beta]\}$

set $CPO(p^*) = (CPO(p^1), p^1)$

set $R_{p^*+1} = \{p^* \cap \{p \in R_{p^*} : F(p) + C(x_{(p+1):p^*}) + K < F(p^*)\}\}$

Record changepoints in CPO

Initialise: $m = \text{length of CPO}$, $k = \text{number of trials}$, $RS = \text{poisson spike train generated from ramp } X(S:E)$, $SS = \text{stepping simulation generated from ramp } X(S:E)$ (see 4.1), startpoint S , $E = \text{endpoint } E$

For: $cp = 1, \dots, m$

Let $S = CPO(cp)$

Let $E = CPO(cp + 1)$

Let $y = \{r : r = X(S) + [\frac{X(E)-X(S)}{E-S}]t, t \in \{0, (E-S)\}\}$

Let

$$f(x) = \begin{cases} x & \text{if } 0 < x < 1 \\ 1 & \text{if } x > 1 \end{cases}$$

Let $Q = \frac{(x(S)-x(E))^2}{\sigma_{x_a:b}^2 - y_{a:b}} \times f(E-S)$

If $X(E) - X(S) > A$ and $E - S > L$ and $Q > 50$:

For: $i = 1, \dots, k$:

Calculate $L = \prod_i \frac{p(V)^{x_a}}{x_a!} \exp^{-\Delta p(V)}$ (see 4.1)

Calculate $AIC = 2k - 2 \ln(L)$

Output: Record AIC values in AICO(i) for ramp simulation, step simulation and data

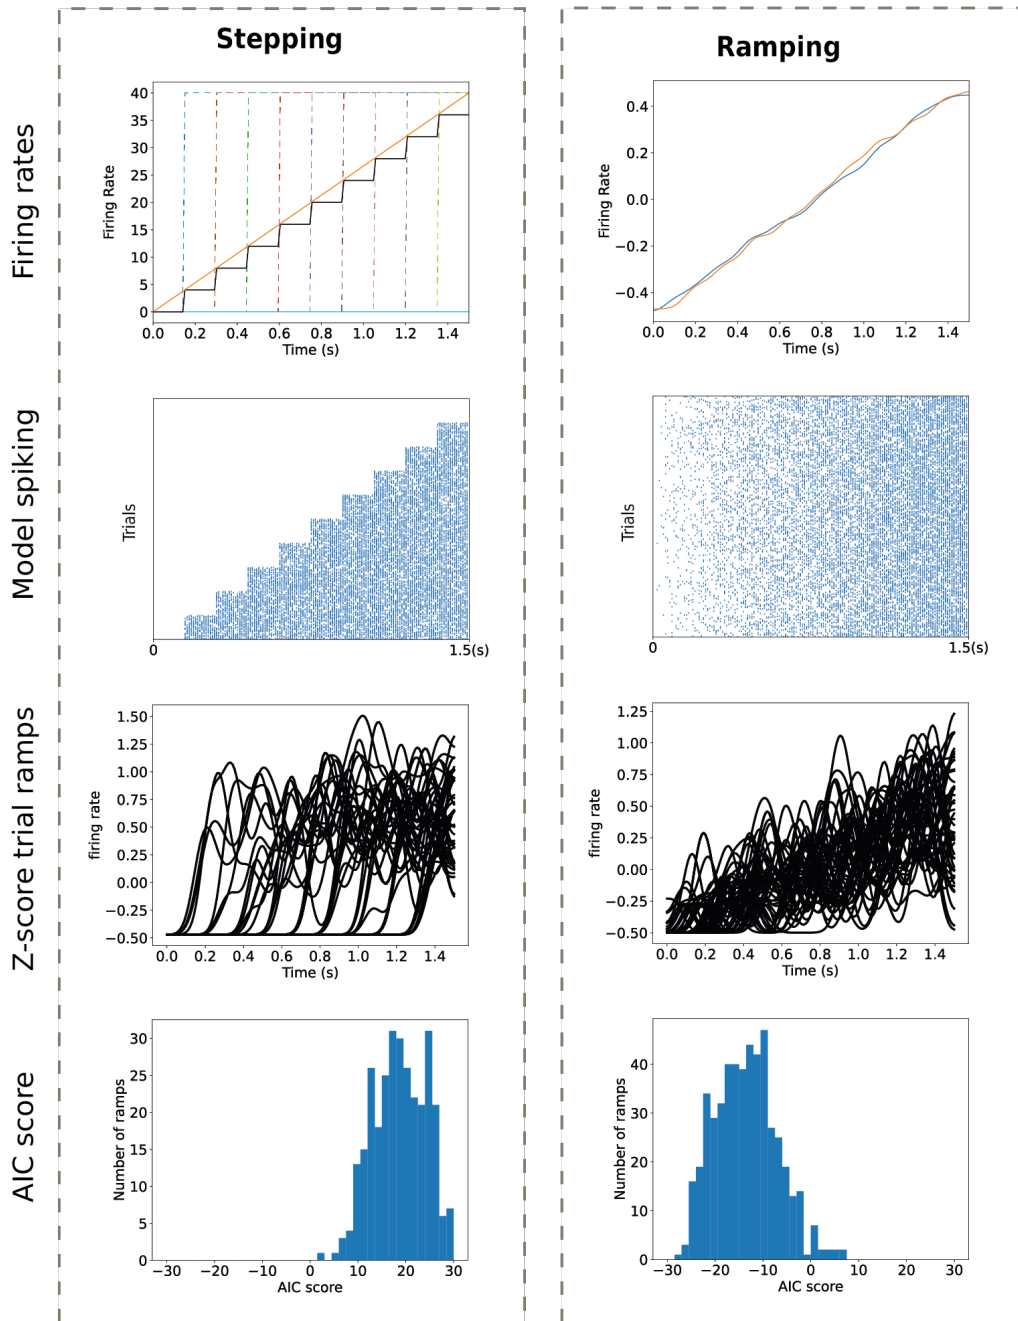


Figure 4.2: Single trial models: building ramping and stepping models. Orange line on top two panels represent the original ramp. Blue line on the top right panel represents the average firing rate of all the simulated trials. Black line on the top left panel represents the average firing rate of all the simulated steps. Dotted lines represent the firing rates for each step. Here, we show a simple example with equal-sized steps. Poisson spiking models are then generated using the ramp or step model (second row rasters). Next, the Z-score trial ramps are calculated for each trial (third row). AIC score histograms show whether the single trials are more step-like (rightward bias, see bottom left) or ramp-like (leftward bias, see bottom right).

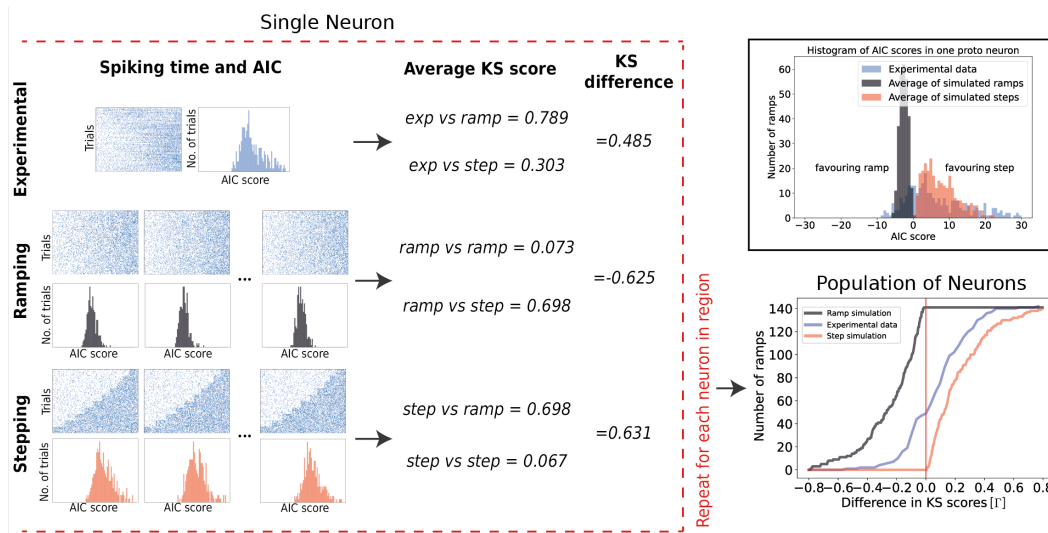


Figure 4.3: Single trial baseline calculation process. First, the AIC difference scores of the firing rates for a single neuron over all trials were plotted. Here we show an example for a GP neuron (see top left of the figure). Using this neuron, we built Poisson firing rates that would produce the experimental ramp if the data came from either a stepping or a ramping model. AIC difference scores were calculated for each trial (see histograms below rasters to left of panel). This was repeated to generate multiple stepping and ramping Poisson spiking and corresponding AIC difference score histograms. In this example, we calculated ten generated models for each of stepping and ramping. Next, the AIC difference score distributions were compared by calculating the KS score between them. Two comparisons were made for each of the experimental data and the stepping and ramping models. The experimental data were compared first with the ramp AIC and then with the step AIC (see top row). For the step and ramping models, each AIC histogram was compared with all others generated from the same model (e.g. step is compared with all other step models AIC, and ramp with all other ramps). Then, the AIC histograms were compared to those from the other model so that step and ramp were compared, and vice versa. For each comparison, an average KS score was calculated and compared with the other average KS scores calculated for each model. For example, see middle row, the KS difference is calculated between the average score from the ramp vs ramp and ramp vs step model. The top right inset shows the average histogram for one neuron with AIC difference scores for each model and experimental data. The KS scores would indicate that the neuron favours the step model. This whole process was then repeated for each neuron in a region where ramp and step models were built and compared. This resulted in three KS difference scores for each neuron. The cumulative distribution of these can be seen in the bottom right. The black line shows the ramp model, the red step model, and the experimental data is the central blue curve. If the KS score of a neuron for that region falls to the left of the red line, it likely favours the ramp model and, to the right, the step model. The area between the red central dashed lines demonstrates the uncertainty zone, where it is unclear if the experimental data was more likely generated from the ramp or step model.

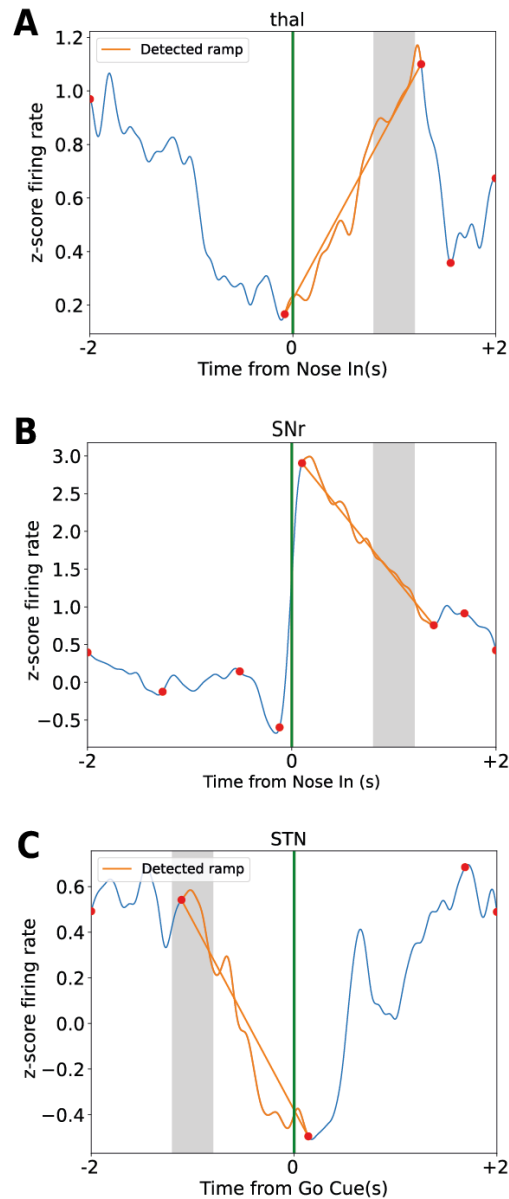


Figure 4.4: Ramp detection with PELT. Examples of ramp detection using the PELT algorithm for the thalamus (panel A), SNr (panel B) and STN (panel C). Z-scored firing rates are calculated between -2(s) and +(2s) around Nose in or the Go cue (Green line) and smoothed with a Gaussian filter. Grey box indicates Go cue in panel A and B, and Nose in in C. Start points are highlighted in red. Where a ramp is detected between the startpoints the firing rate for this piecewise segment and fitted linear ramp are shown in orange. Panel A shows a positive ramp, and panel B,C, show negative ramps.

and potentially receive a food reward. Therefore, they had to exert action restraint to prevent any premature movement out of the port.

To identify and examine ramping in the average firing rate of each neuron, the neural activity could be aligned to the start or endpoint of the hold period (i.e. the Nose in or Go cue event, respectively). As the duration of the hold period was randomized (in order to prevent the animals from being able to predict the onset of the Go cue), the alignment of the neural activity affected the interpretation. For example, transient response to the Go cue could create a spurious ramp starting towards the end of the hold period when aligned to the Nose in event (see fig 4.5). In both examples shown in fig 4.5, a ramp was detected in the Nose in alignment, but not the Go cue alignment. When we align the firing rates to Nose in, individual trials have different Go cue timings, and strong responses to the Go cue could be staggered such that a spurious ramp is formed when aligned to Nose in. Similarly, ramps aligned to the Go cue could result from strong responses to Nose in. Therefore we detected ramps for both alignments separately and required that a ramp was present in both alignments to avoid spurious ramps entering the analysis.

Overall, all basal ganglia subregions (and the thalamus) contained a significant number of neurons with a ramping firing rate during the hold period (8-13.9%; see Table 4.1). Noticeably, the highest proportion of ramping neurons was present in thal (13.9%). The proportion of ramps that were due to increasing and decreasing firing rate (i.e. positive or negative ramps) were similar in all regions. We conclude that a subset of neurons in all basal ganglia subregions (as well as the thalamus) exhibited ramping activity while animals were waiting for a Go cue. However, the proportion of ramping neurons was lower than in some previous reports (Donnelly et al., 2015; Mirzaei et al., 2017), in line with a more conservative detection algorithm.

4.3.2 Ramp properties differ across the basal ganglia.

In the next step, we then examined the detected ramps in each subregion in more detail. For each subregion we visualized the average ramp (separately for positive and negative ramps) of each neuron with a line indicating the ramp's start and endpoint relative to the Nose in and Go cue events (fig. 4.6). For ramping activity that merely occurs randomly and independent of the task events, a null hypothesis

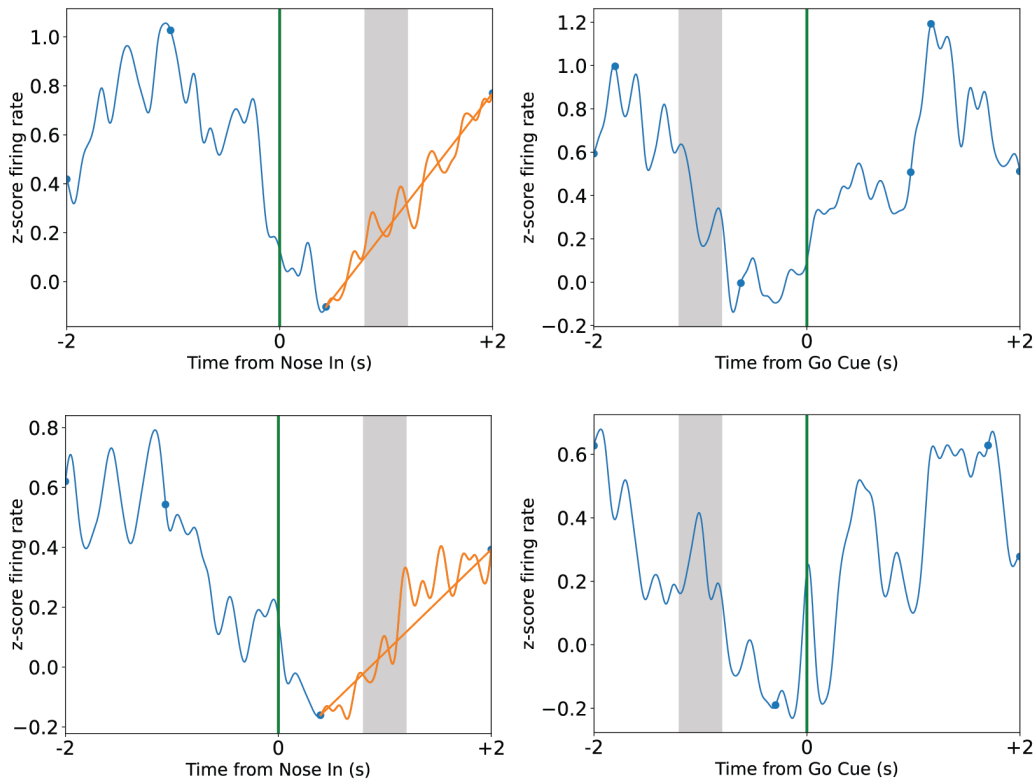


Figure 4.5: Sharp transient activity after Go cue can result in ramp-like activity in Nose in aligned scenario. Each row shows the firing rate of an STN neuron aligned to both Nose in and Go cue. Only detected ramps that start within 0.5 seconds of the Nose in alignment and between -0.5 seconds to -1.5 seconds before the Go cue in Go cue alignment were highlighted. Our approach of observing ramps in both event alignments ensured that spurious ramps like this were avoided.

Region	# Neurons	# Ramping neurons during hold	# Positive ramps	# Negative ramps
Str	887	70 (7.9%)	42 (4.7%)	28 (3.2%)
GP	515	58 (11.3%)	32 (6.2%)	26 (5.1%)
SNr	207	27 (13.0%)	18 (8.7%)	9 (4.4%)
STN	258	30 (11.6%)	13 (5.0%)	17 (6.6%)
Thal	345	48 (13.9%)	20 (5.8%)	28 (8.1%)

Table 4.1: **Ramping activity in basal ganglia neurons is common while animals wait during the hold period.** For each region, ramps during the hold period were detected using the PELT algorithm, with the data being aligned either to the Nose in or the Go cue event. We only included ramps that were detected in both alignments. The total number of recorded neurons is given (column 3) together with the number of neurons with an identified ramp during the hold period (column 4). Identified ramps were then further split up into ramps with an increasing firing rate (positive ramps) or decreasing firing rate (negative ramps).

is that the ramp onset time points are uniformly distributed in the time window. To account for the varying duration of the hold period, we then determined start and endpoints of all ramps for Nose in and Go cue alignment respectively for more accurate measurement of ramp properties, as well as their duration and amplitude.

For the analysis of the ramp properties, one main question was whether the ramps in different regions had overall the same properties. We performed ANOVAs on various ramp properties (see Table 4.2). Length, amplitude, and slopes are calculated from the difference between the Go cue aligned endpoint and Nose in aligned startpoint. We found significant differences in amplitude for negative ramps after multiple comparisons correction (see Table 4.2). In addition, we found that the mean value of amplitude for negative ramps was significantly different between SNr and GP ($p = .0014$, t-test for multiple comparisons), SNr and thal ($p = .0001$, t-test for multiple comparisons). In general, SNr had negative ramps with much lower amplitudes (see Table 4.3). No significant differences in amplitude were found for positive ramps (see Table 4.2).

We found no significant differences between regions in the startpoints for both negative and positive ramps. Most regions have a mean near the Nose in event (see Table 4.3), although this might result from our constraints on startpoints.

We found no significant differences between regions in the endpoints for positive and negative ramps. Most regions have a mean near the Go cue event (see Table 4.3). However, standard deviations for positive ramps' endpoints are smaller and closer to 1 second (the go cue) than those in negative ramps. This suggests that negative ramps tend to end later while positive ramps tend to end near the go cue.

It should be noted that standard deviations are higher in endpoints than startpoints as they have less of a constraint. Startpoints are filtered only within .5 seconds of the Nose in event, meaning it has a range of possible values up within 1 second. However, endpoints could end anywhere up to 2 seconds after the Go cue, which means the range of possible values could be up to 3 seconds.

We found no significant differences between regions in the lengths for negative and positive ramps. Most regions have a mean length between 1- 1.4 seconds (see Table 4.3).

In general, SNr had strong negative ramps; these ramps were also of the highest mean length (see 4.8, although no significant differences in length were found. This is reflected in SNr negative ramps ending and starting later than every other region on average (see Table 4.3). They also had the most negative slopes at - 2.2 z-score per second. These negative ramps also had a far more concentrated distribution of startpoints (see fig. 4.7, suggesting that they could be a similar subset of ramps that play the same role in the behaviour.

STN has positive ramps that start shortly before Nose in fig. 4.7, similar to the startpoints of the negative ramps in the SNr; they are also concentrated and have a low standard deviation (see Table 4.3). These ramps were, on average last to start when aligned to Nose in and have the strongest slope (1.99 z-score per second).

To test whether ramps that started later in the hold period ended later, we looked at correlations between startpoints and endpoints of various regions. No regions were found to have significant correlations after corrections for multiple comparisons.

To better understand the activity of these ramping neurons throughout the task, we calculated the population average of these ramping neurons for each alignment fig. 4.9. We did not only look at the segments that were ramping but the entire firing rate around the trial event. In general, and in all regions, while positive ramps tend to end at the Go cue, negative ramps tend not to end later (e.g. in the thalamus and Str in fig. 4.9).

We can clearly see that the firing rates with SNr negative ramps had a sharp increase in activity near the Nose in event. However, the ramp gradually decreased over the hold period until shortly after the Go cue, when there was a small increase in activity. This supports the idea that these negative ramps are responsible for action restraint in the hold period.

STN and SNr both have stronger z-score positive ramps 4.9. These ramps also more clearly reflect the reaction to Nose in, unlike other regions where population averages seem to start earlier.

Overall, these strong positive STN, positive SNr and negative SNr ramps are suggestive that these regions are involved in the hold period and both reactive inhibition (such as action cancelling in the stop signal task (Schmidt et al., 2013)) and proactive inhibition, such as action restraint in hold period.

Ramp direction	Startpoint (F, p)	Endpoint (F,p)	Length (F,p)	Amplitude (F,p)	Slope (F,p)
Positive	(.89, .47)	(.43, .78)	(1.04, .38)	(1.1, .35)	(1.9, .1)
Negative	(2.6, .04)	(1.4, .24)	(.73, .57)	(7.0, $p < .0001$)**	(3.78, .0067)*

Table 4.2: **ANOVA shows significant differences in amplitudes of negative ramps in the hold period.** Table of ANOVA results for each ramp characteristic. F value is displayed first, followed by the p-value. P-values are before multiple comparisons. * represents that the ANOVA has $p < .05$. ** represents that the ANOVA has $p < .05$ after Bonferroni correction. Start points were taken from Nose in aligned trials. End points were taken from go cue-aligned trials. Length, amplitude, and slopes are calculated from the endpoint and startpoint.

Region	Positive/ Negative	Start (s)		End (s)		Length (s)		Amplitude (z-score)		Slope (z-score/s)	
		Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
Str	P	-0.13	0.24	0.93	0.32	1.06	0.30	1.33	0.63	1.34	0.82
GP	P	-0.12	0.21	1.02	0.42	1.14	0.35	1.25	0.75	1.24	0.99
SNr	P	-0.08	0.23	0.92	0.33	0.99	0.38	1.65	1.24	1.93	1.81
STN	P	-0.01	0.12	0.95	0.39	0.96	0.33	1.54	1.00	1.99	2.00
thal	P	-0.10	0.19	0.96	0.18	1.05	0.27	1.24	0.38	1.24	0.45
Str	N	-0.13	0.18	0.92	0.53	1.05	0.45	-1.28	0.69	-1.38	0.88
GP	N	-0.11	0.22	1.12	0.48	1.23	0.50	-1.16	0.35	-1.07	0.47
SNr	N	0.11	0.18	1.34	0.43	1.23	0.34	-2.44	1.88	-2.12	1.73
STN	N	-0.07	0.26	1.03	0.58	1.10	0.40	-1.11	0.45	-1.08	0.42
thal	N	-0.05	0.21	1.19	0.60	1.24	0.53	-1.08	0.43	-1.06	0.72

Table 4.3: SNr shows strong negative ramps that start later than the rest of the regions. Means and standard deviation were calculated for each ramp property for all regions.

4.3.3 Ramp characteristics do not reflect single trial behaviour

Ramps in the average firing rate of a neuron can be due to ramping activity in single trials, but, importantly, also due to step-like changes at different time points, which then appear as a ramp only in the trial average. To determine whether a single trial is more likely to have a ramping or step-like firing rate change, we compared the experimental data with simulations based on spike trains with either ramping or step-like firing rate changes (see Methods). Our method then yielded an overall score Γ_i for each neuron i , with positive values indicating that the ramp was due to steps in single trials and negative values indicating that the ramp was due to ramps also in single trials (see top row of fig. 4.10). More examples of single trial AIC distributions and raster plots can be found in Appendix B.6.

To determine whether a given subregion shows single trial ramping or stepping, we examined the resulting cumulative distributions of Γ across neurons from the same subregion. We obtained three distributions of Γ : Γ^{exp} from the experimental data, Γ^{ramp} from the ramp simulation and Γ^{step} from the step simulation (fig. 4.11, 4.12).

We found that most regions have many unclear neurons (defined by a less than .1 difference in KS scores) that are neither stepping nor ramping (see fig. 4.11). Also, most regions have slightly more stepping neurons than ramping neurons. Noticeably, the striatum has the highest number of stepping neurons and SNr negative ramps seem to be more ramp-like in general.

In the positive ramps of the SNr (see fig. 4.11), the Go cue aligned trials of SNr are more step-like, while the Nose in aligned trials of the SNr are ramp-like. This is mainly because SNr ramps have a very strong increase in activity at Nose in (see fig 4.9). If this strong increase in activity is within the parameters of our definition of a ramp, it will be captured by our ramp detection algorithm. This ramp will have startpoints that are not influenced by the variable delay in Go cue. Their endpoints are similarly not affected as they end before the Go cue. However, when aligned to the go cue, this activity will present as a longer ramp, as though several step-like ramps are averaged across trials. The longer length of the detected Go cue aligned ramp will then capture the plateaus or dips at the end of the strong ramp in the single trials. Thus, we find that for the SNr, there

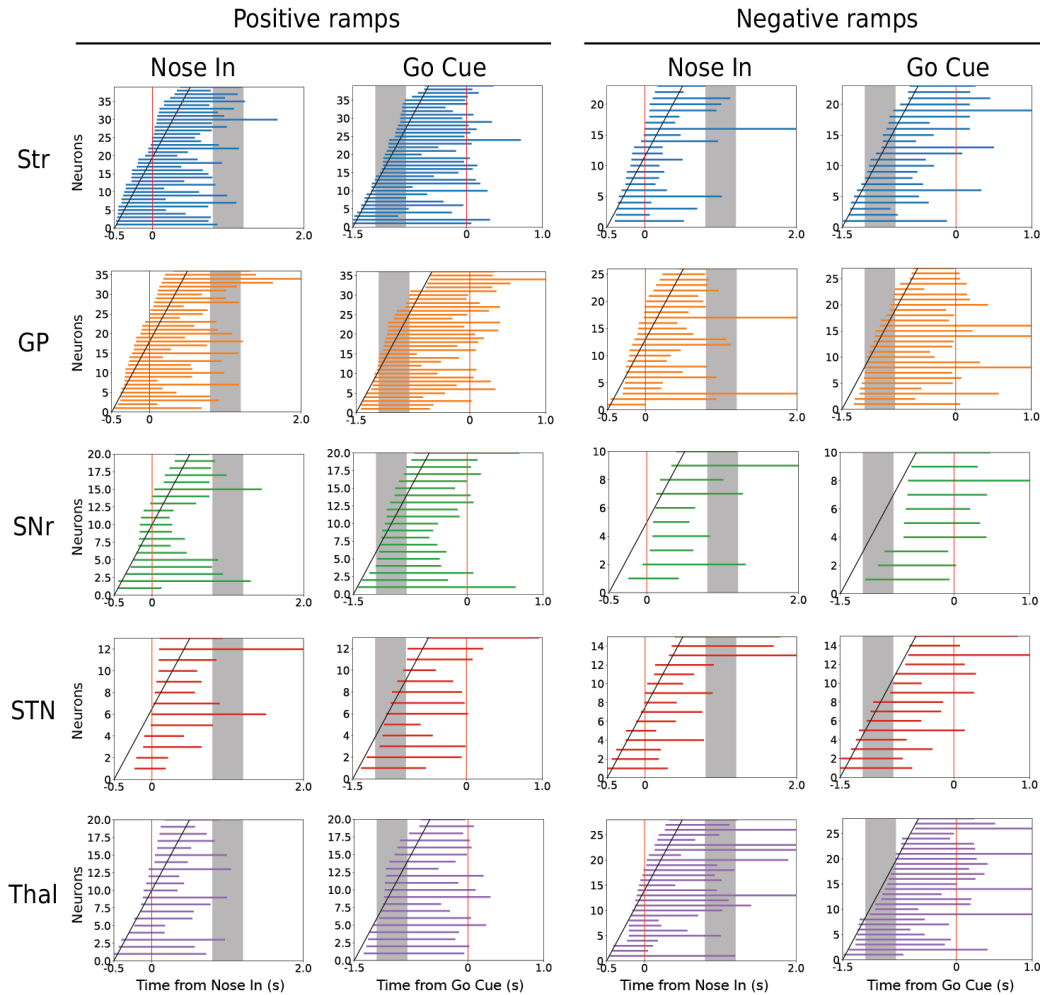


Figure 4.6: Single neuron ramps cover the hold period in all basal ganglia subregions. The four columns show positive and negative ramps separately for alignment to the Nose in and the Go cue event, respectively. Each row shows a different basal ganglia subregion (and thalamus). In each panel, the coloured line covers the ramp of a single neuron's ramp from its startpoint to its endpoint. Within each subregion, the ramps were ordered by their starting point. The grey box indicates the time range in which the Go cue could occur (for Nose in alignment) and the time range in which the Nose in event could occur (for Go cue alignment). The diagonal black lines mark the ramp starting points for hypothetical uniform starting point distributions.

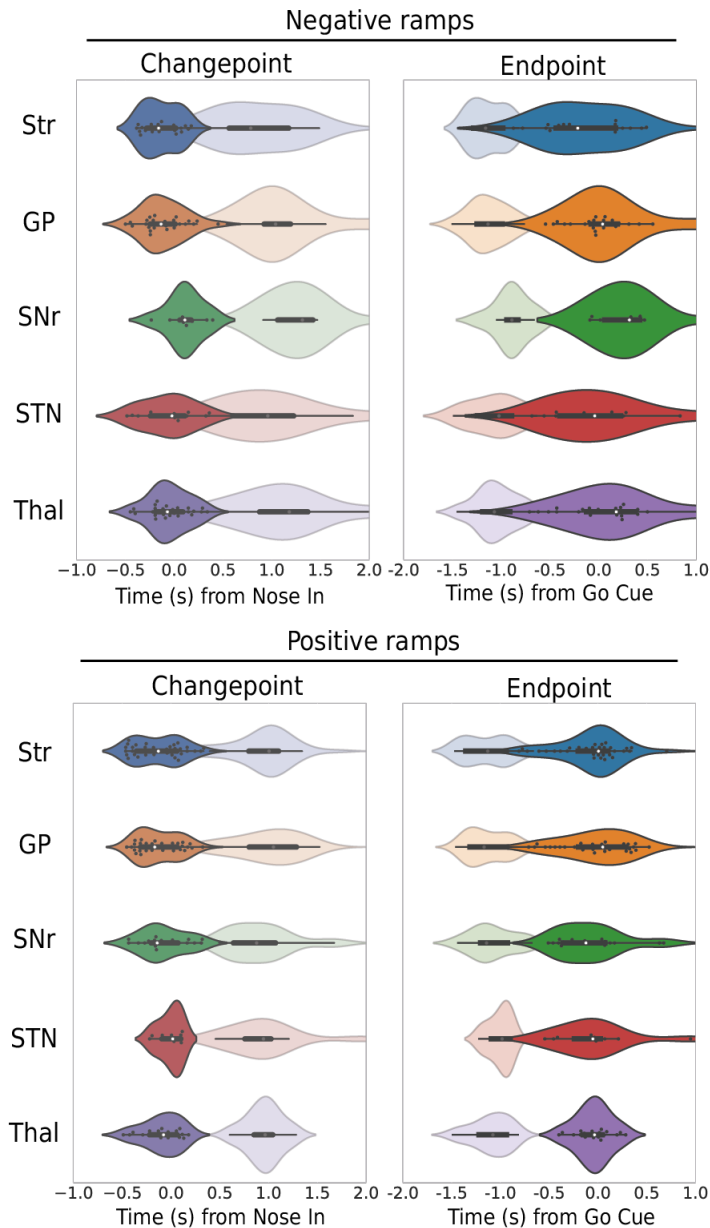


Figure 4.7: STN positive ramps start close to the go cue. Violin plots of startpoints and endpoints of detected ramps. The startpoints and endpoints of detected negative and positive ramps. Start points were taken from the Nose in alignment, while endpoints were taken from the Go cue alignment. Faint coloured violin shows the endpoint distribution in the startpoint violin plot and vice versa. Most ramps start at a mean of around one second before the Go cue, on average. This is the mean of where the hold period usually starts. Similarly, most endpoints end around the Go cue.

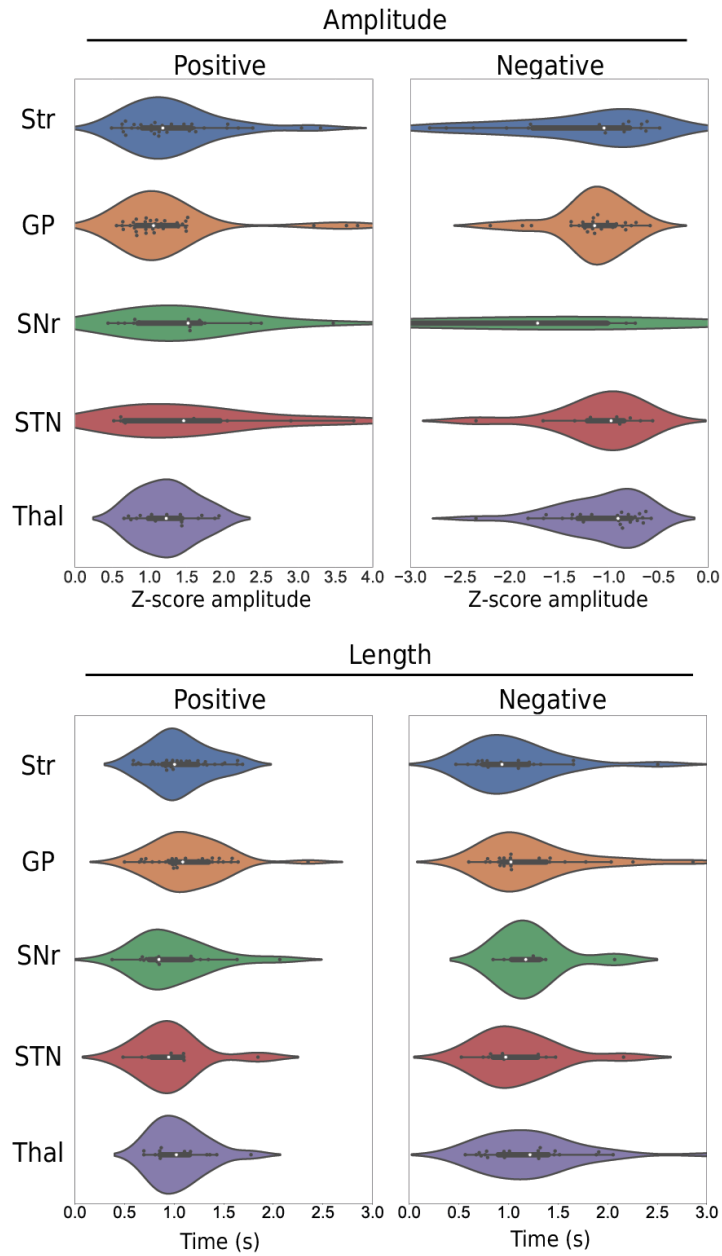


Figure 4.8: SNr has strong negative ramps (amplitudes are much lower) compared to other regions. Lengths are very similar across regions. Violin plots of amplitude and length of the detected ramps for both positive and negative ramps. The length and amplitude are both derived from deducting the respective values of the Go cue aligned endpoint from the Nose in aligned startpoint.

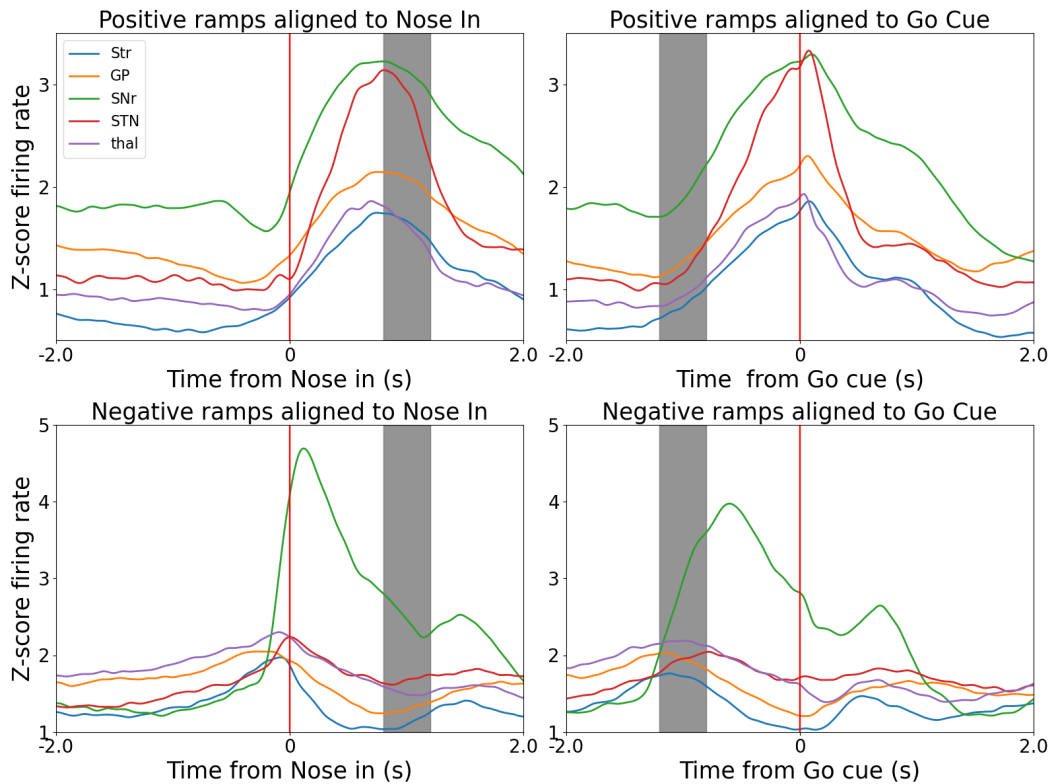


Figure 4.9: Population average of ramping neurons in each region show ramps end shortly after movement onset. We averaged the firing rate across all neurons of a population. STN and SNr show strong population average positive ramps. SNr shows strong population average negative ramps.

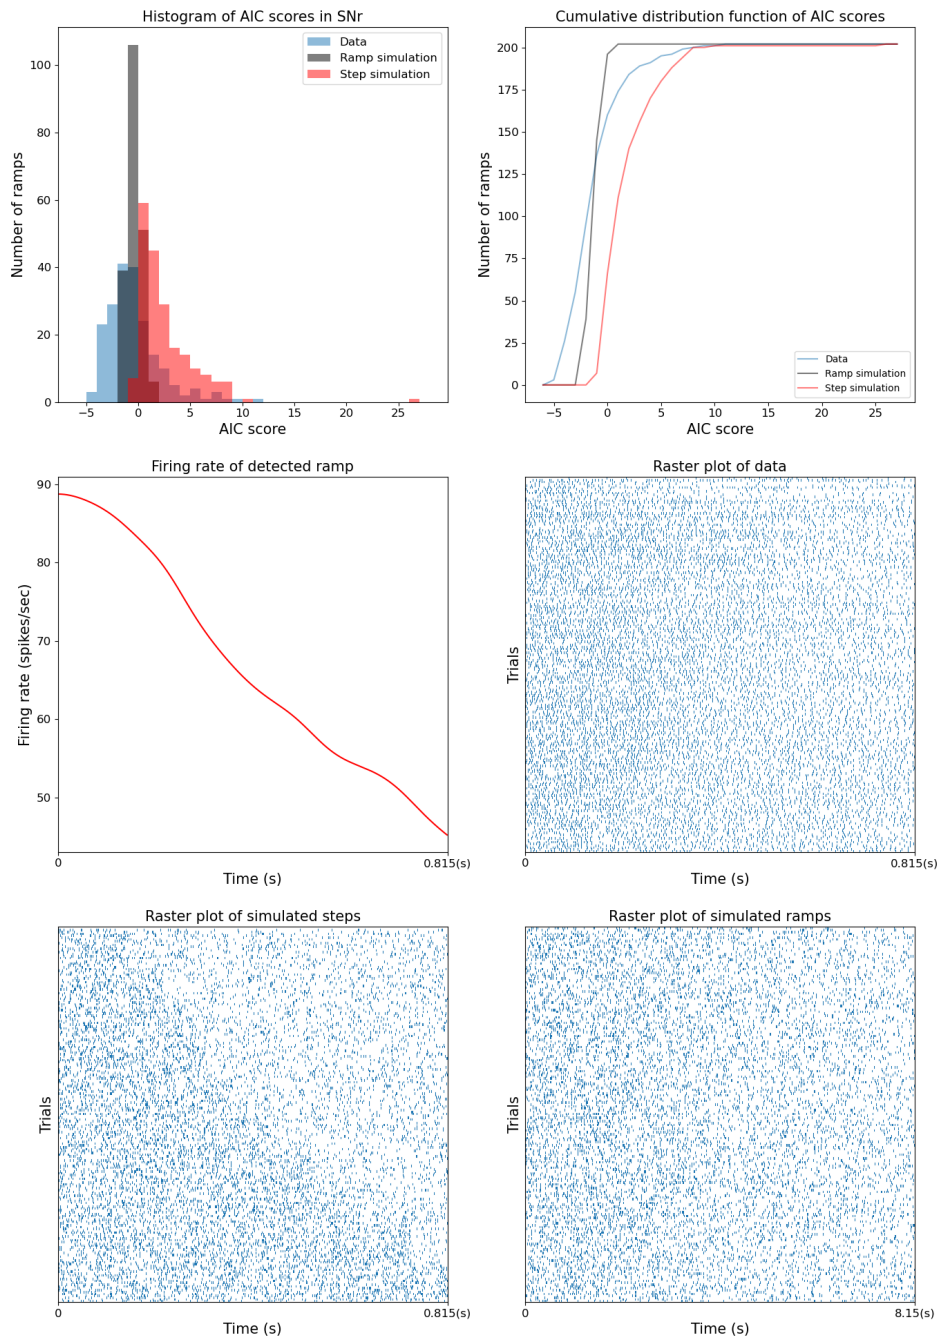


Figure 4.10: Detected ramp in SNr has ramping single-trial firing rates. Top row shows AIC score distribution on the left and the cumulative distribution of AIC scores on the right. Middle row shows the firing rate of the detected ramp on the left with the raster plot of the trials on the right. Bottom row shows raster plots of simulated steps and ramps.

are differences in Go cue and Nose in single trial activity despite comparing the same data.

AIC scores in experimental data can be higher or lower than in simulations. For example, the poisson firing rates are built out of the average ramp, with a baseline firing rate at the start. However, ramps can start at 0 firing rate for some trials to construct a 'perfect ramp', while the baseline firing rate is built out of other trials. Despite this, populations of neuronal firing rate have lower Γ values than simulations (see fig. 4.11, 4.12).

We might expect that ramp shape and characteristics vary between stepping neurons and ramping neurons. For example, single-trial stepping neurons and single-trial ramping neurons could form two different subtypes of ramps, which is reflected in the trial-averaged firing rates. To find that out, we calculate the correlation score between Γ^{exp} and the various ramp characteristics, such as firing rate, startpoint, endpoint, length and amplitude (see Table 4.4). However, for most regions and ramp characteristics, there is no significant correlation between them, suggesting that we are unable to accurately capture how ramping single neurons respond in trials from their average firing rates.

4.4 Discussion

We adapted methods from the changepoint and the wind ramp detection literature to detect and characterise ramps in the basal ganglia. More specifically, we used PELT to detect changepoints in the time series of a firing rate and generated piecewise segments from those changepoints. Finally, we require thresholds for minimum amplitude and length to detect a ramp, similar to many papers in the wind ramp literature. Our method allows us to accurately characterise when a ramp starts, ends, and its amplitude, unlike most past analyses, which identify this qualitatively. This allowed us to more objectively identify and characterise individual ramps that start at the hold period. Furthermore, the definition of a ramp in neuroscience is lacking, with some papers looking at ramps as low as .3 seconds (see Catanese and Jaeger (2021), and others as long as 12 seconds see (Emmons et al., 2017)). Varying ramp properties could signify different types of function.

In the 'hold your horses' models of the STN (Frank, 2006), ramps in the hold

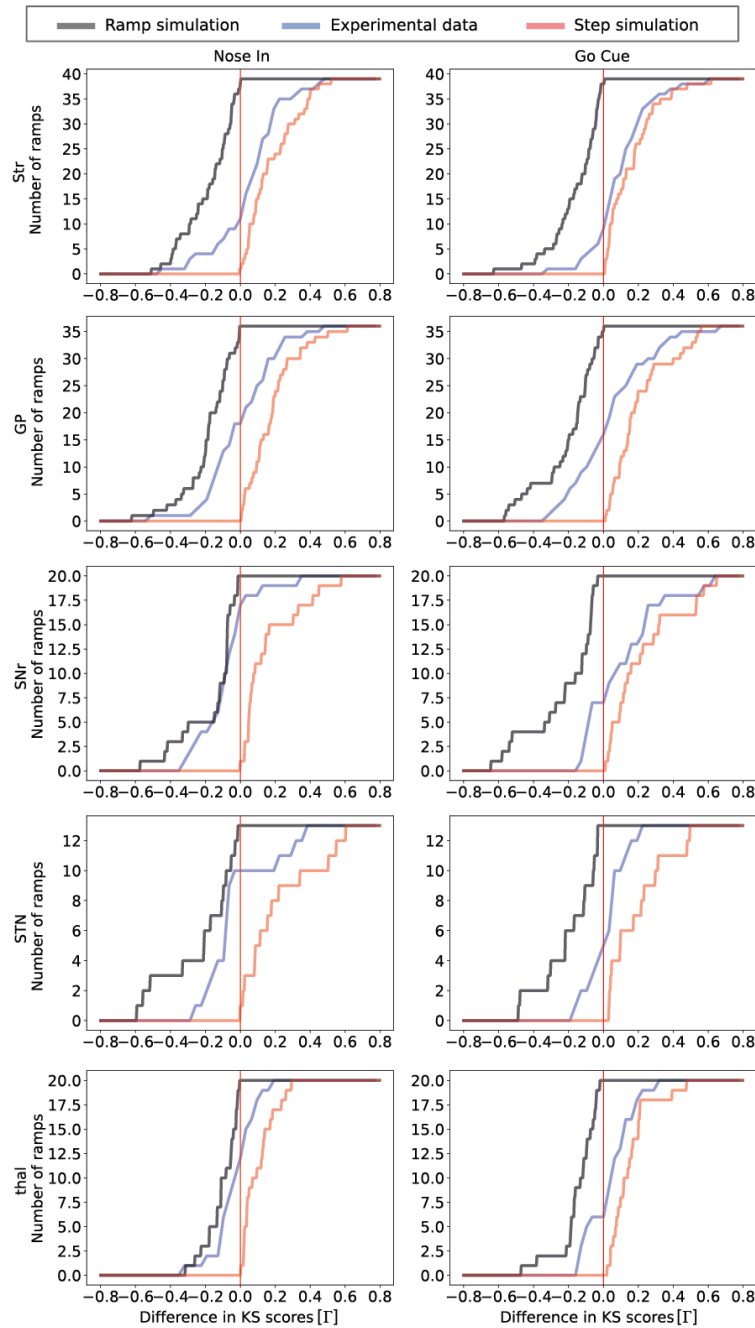


Figure 4.11: Positive ramps in all regions are a mix of stepping and ramping single-trial activity. Cumulative plots of single trial simulations for positive ramps. Cumulative number of ramps in each region for each difference in KS score, aligned to Go cue or Nose in. When the experimental data is closer to the red line, this indicates more stepping-like neurons. Whereas when it is closer to the black line, it indicates more ramping-like neurons.

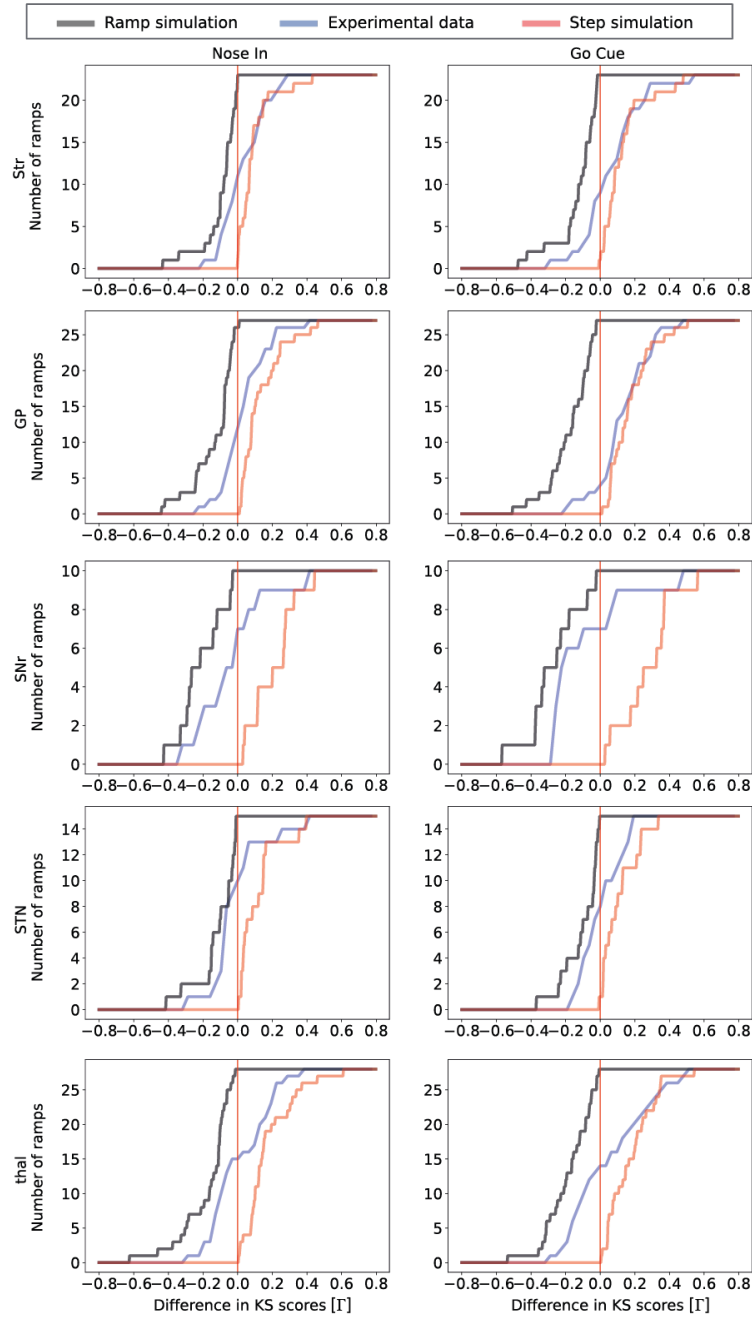


Figure 4.12: Negative ramps in SNr and STN are mostly ramping in single trials. Cumulative plots of single trial simulations of negative ramps. Cumulative number of ramps in each region for each difference in KS score, aligned to Go cue or Nose in. When the experimental data is closer to the red line, this indicates more stepping-like neurons. Whereas when it is closer to the black line, it indicates more ramping-like neurons.

Region (No. neurons)	Alignment	No. positive ramps	Stepping neurons (positive)	Ramping neurons (positive)	Correlation Coefficient				
					Quality	Length	Start FR	Amplitude	FR
Str (887)	Nose in	39	20(51.3%)	5(12.8%)	-0.26	0.26	0.03	-0.22	-0.00
Str (887)	Go cue	39	20(51.3%)	2(5.13%)	-0.22	0.19	0.23	0.36*	0.32
GP (515)	Nose in	36	13(36.1%)	9(25.0%)	-0.33	-0.22	-0.13	0.03	0.01
GP (515)	Go cue	36	13(36.1%)	9(25.0%)	-0.10	-0.13	-0.15	0.03	-0.05
SNr (207)	Nose in	20	2(10.0%)	5(25.0%)	0.01	0.37	0.13	-0.12	-0.02
SNr (207)	Go cue	20	10(50.0%)	0(0%)	0.12*	-0.27	-0.07	0.14	0.07
STN (258)	Nose in	13	3(23.1%)	3(23.1%)	0.47	-0.01	-0.01	0.44	0.27
STN (258)	Go cue	13	3(23.1%)	1(7.69%)	0.56	-0.30	-0.31	-0.03	-0.14
thal (345)	Nose in	20	4(20.0%)	2(10.0%)	0.12	-0.35	-0.16	-0.22	-0.19
thal (345)	Go cue	20	7(35.0%)	3(15.0%)	-0.36	-0.29	-0.65*	-0.62*	-0.62*

Table 4-4: **Stepping neurons are more prevalent than ramping neurons in most basal ganglia regions.** Detected ramping and stepping neurons from the single-trial analysis for positive ramps. Nose in aligned ramps have more stepping neurons in most regions. Correlation coefficients for each of the ramp characteristics with Γ_i . The higher the Γ_i , the more stepping the neurons are. FR refers to firing rate. * represents that the correlation has $p < .05$. ** represents that the correlation has $p < .05$ after Bonferroni correction.

Region (No. neurons)	Alignment	No. negative ramps	Stepping neurons (negative)	Ramping neurons (negative)	Correlation Coefficient				
					Quality	Length	Start FR Amplitude FR		
Str (887)	Nose in	23	9(39.1%)	1(4.4%)	-12	-.20	-.39	-.53*	-0.46*
Str (887)	Go cue	23	11(47.8%)	2(8.7%)	.20	-.37	-.16	-.37	-0.25
GP (515)	Nose in	27	8(29.6%)	2(7.4%)	-.39*	.34	.26	.26	0.25
GP (515)	Go cue	27	19(70.4%)	2(7.4%)	-.10	.05	.09	.05	0.08
SNr (207)	Nose in	10	2(20.0%)	3(30.0%)	-.36	.06	-.45	-.58	-0.53
SNr (207)	Go cue	10	2(20.0%)	6(60.0%)	-.61	.07	-.36	-.46	-0.37
STN (258)	Nose in	15	2(13.3%)	2(13.3%)	-.07	.41	.34	.33	0.37
STN (258)	Go cue	15	5(33.3%)	2(13.3%)	-.47	-.30	.33	.38	0.35
thal (345)	Nose in	28	12(42.9%)	7(25.0%)	.19	-.46*	-.42	-.49*	-0.46
thal (345)	Go cue	28	12(42.9%)	7(25.0%)	.30	-.30	-.50*	-.53*	-0.53*

Table 4-5: **SNr negative ramps are reflected in single-trial activity too.** Detected ramping and stepping neurons from the single-trial analysis for negative ramps. Nose in aligned ramps have more ramping neurons in SNr, STN, thal, GP, and more stepping neurons in Str. Go cue aligned ramps have more stepping neurons than Nose in alignment neurons. Correlation coefficients for each of the ramp characteristics with Γ_i . The higher the Γ_i , the more stepping the neurons are. * represents that the correlation has $p < .05$. ** represents that the correlation has $p < .05$ after Bonferroni correction.

period might be involved in preventing premature movements before the Cue to receive a reward. Many of these STN ramps were similarly found in (Mirzaei et al., 2017). However, despite using the same dataset, we found fewer ramping neurons compared to (Mirzaei et al., 2017). This is generally because our approach to ramp detection is more conservative than them, in particular, we only selected ramps in both Nose in aligned and go cue aligned trial-averages. This selection removed many ramps that had spurious ramps in the STN, where the Nose in aligned ramps were due to transients in the Go cue. Furthermore, we also quantified the start of these ramps and found that many of these STN ramps start at the onset of Nose in, providing further support for the 'hold your horses' model of the STN. Still, we found only 11.6% of ramping neurons in the STN compared to the 77% found in Mirzaei et al. (2017). Our single trial analysis of these ramps suggests that it is unclear whether they are stepping or ramping and is likely a mixture of both. This contradicts the hypothesis that ramping in the STN is preparatory. More complex firing rates can also result in ramping on the trial-averaged firing rates, which could be investigated in future research.

We found strong negative SNr ramps. Surprisingly, these ramps lasted slightly longer than the go cue on average and started slightly after the Nose in event. This could be more in line with the movement initiation of the rat rather than the Go cue itself. Downramps in GPi (Similar to the SNr in the rat) were found in primates in the hold period before reaching movement (Thura and Cisek, 2017). This is in line with what we found here in the SNr. Overall, negative ramps in downstream structures of the basal ganglia support findings of positive thalamic ramps before movement (Catanese and Jaeger, 2021). They also found that single trial units tend to have stepping firing rates but sum up to a ramping firing rate. They also found that their linear decoder could better determine trial outcomes when they used more neurons. This is in line with our analysis where single trials are mostly stepping and discrete and not ramping. Future work could correlate these single-trial properties to behaviour, such as reaction time.

Our selection of a minimum z-score of .5 is conservative. In general, our conservative algorithm suggests that many studies looking at ramps out there should take caution when analysing ramps (see appendix B.7 upper right and bottom left panel for some examples of rejected ramps just under .5 z-score). Many ramps have been found to be noisy and have a very small increase in the z-score firing

rate. Our initial ramp detection algorithm, without the quality measure, captured many of these ramps (analysis not shown). Compared to other studies, our selection of a .35 second minimum ramp could also be conservative (see Catanese and Jaeger (2021) for 1.3 second ramps.) Furthermore, they investigated firing rate ramps and the threshold model of action release, suggesting that action gets initiated upon threshold release. This requires a study of firing rate ramps without the z-score normalization.

While there is a debate between stepping and ramping single trials in the LIP (Zhao and Kording, 2018; Latimer et al., 2015; Zoltowski et al., 2019; Kiani and Shadlen, 2009; Kiani et al., 2008), few studies have looked at stepping and ramping single trials in the basal ganglia. More specifically, the LIP studies involved how evidence was accumulated in ramping trials but not in how ramping neurons could be involved in inhibitory control. Here, looking at various basal ganglia regions, we show that while the frequency of ramping and stepping neurons varies across regions, striatal neurons, in general, tend to have more stepping neurons than ramping neurons. While stepping neurons in the hold period could be related to action suppression, it is unlikely that these stepping neurons result from anticipation of an action; for example, the animal thinks there is an increasing chance that the Go cue is about to go on. We also find that, in most regions, there is no significant correlation between ramp characteristics and whether the single trials are ramping or not. This suggests that future studies should carefully consider how the single trials are firing when investigating ramping neurons. This is important for tasks that look at average ramps preceding the event alignment, especially when the events vary in time (for example, the Go cue here could vary between .8s and 1.2s) as individual steps in firing rate that vary in start time could add up to trial-averaged ramps.

We have only used correct Go trials in our analysis above as we were observing inhibitory behaviour in the hold period. However, failed Go and incorrect Go trials are much fewer and the sample size is too small to find enough ramping neurons. Incorrect Go is where the rat goes into the wrong port (left when they should be going right, and vice versa). Failed Go is where the rat fails to move out of the centre port in time. With more trials for failed Go, we might observe differences in ramping if inhibitory behaviour requires ramping to expect and time events. For example, ramping behaviour could be disrupted more in failed Go

trials. Furthermore, past studies have shown that ramping could also start earlier in failed trials (Donnelly et al., 2015). Other studies show that ramping activity was higher for premature licks and weaker for trials with no licks (Catanese and Jaeger, 2021). There could also be differences in amplitudes in ramps between failed and correct Go. If movement requires ramping to hit an upper threshold, then a small amplitude could mean that the rat fails to execute the movement.

Our method allows for comparison of different lengths of ramps, and helps to solve the problems that AIC is known to be fickle between two distinct models, we used simulations to help solve that problem.

Chapter 5

Artifacts in spike-phase entrainment

5.1 Chapter overview

In the previous chapter, we observed ramping firing rates within the hold period of a stop-signal task. These inhibitory behaviours have previously been observed in oscillatory field potentials, for example, decreased alpha activity is linked to a decrease in inhibitory control (Knyazev, 2007). Furthermore, Beta oscillations have been found to be related to cue utilization (Leventhal et al., 2012), and gamma oscillations have been found to be related to rewards (Sevlian and Rajagopal, 2013). Finally, the identification of arkys and protos is difficult. Preproenkephalin is the sole marker to identify arkys from protos (Abdi et al., 2015). Thus there is a need to identify these two subpopulations to better study them. The use of entrainment to various frequencies could be a useful method to identify them. Thus, we first investigated whether arky-pallidal and prototypical neurons are entrained to gamma oscillations in the basal ganglia.

When the spike and LFP from the same tetrode were used for the entrainment analysis, there was a bimodal distribution of entrainment in the GP where there existed two peaks, one shortly after the trough of a gamma cycle and one on the down phase of a gamma cycle. Upon further investigation, we found that the arky cells and proto cells entrain to different phases of the gamma wave in awake

rats. However, this entrainment disappeared when we shuffled the LFPs and spikes (each spike train is tested for entrainment against an LFP from a different tetrode). We reasoned that one possibility was that arkys and protos had different waveforms, and typical pre-processing steps of the LFP could lead to spurious entrainment, such as the one we found above.

In previous studies, arky and proto neurons were not found to have waveform differences in vivo (see supplementary figures in Mallet et al. (2016)), but afterhyperpolarization potential of arkys was found to be twice as strong as that of the proto neurons in vitro (see table 1 in Abdi et al. (2015)). The data used in this chapter was also used in Mallet et al. (2016) in the freely moving animals. As the differences between single-unit waveforms are usually explained by a small number of principal components (Fee et al., 1996), we also looked at the coefficients of the first few principal components of arky and proto neuron waveforms. We show that there are waveform differences between arkys and protos and that both the first and second principal components capture the differences between them. In our dataset, we find that both pre and post-spike normalized voltage was, on average lower in arkys than in protos.

Finally, we wanted to test whether these waveform differences resulted in the spurious entrainment we found earlier. Although studies have looked at the influence of spikes on the LFP signal (Waldert et al., 2013) and found some artifactual correlations between LFPs and spikes Zanos et al. (2011), none have investigated the influence of different waveforms on entrainment. Using typical LFP pre-processing steps, we were able to replicate the mean phases of many arkys and protos with just their unfiltered, raw spike-triggered average (STA) waveform.

5.2 Methods

5.2.1 Dataset

We used the same GP neurons from chapter 4 and limited the analysis to correct Go trials. They were recorded using electrodes in vivo while the rat was engaged in a stop signal task. We used the same classification for putative arky and proto-cells as in Mallet et al. (2016) by using the first principal component of slow wave sleep (SWS) firing rate, coefficient of variation (CV) during SWS and SWS rate

change.

5.2.2 LFP processing

For spike phase entrainment, we obtained session-wide LFPs where both putative arkys and putative protos were recorded. We used the same tetrode for the single unit spiking and the LFP. We used a lowpass filter of 500Hz, then downsampled the LFP from 31250 to 496 Hz, and lastly, we bandpass filtered the LFP by the gamma frequency (45-55Hz) and calculated phases with a Hilbert transform. Finally, we obtained the phases where spikes occurred on the session-wide LFP and tested for non-uniformity (Rayleigh test).

5.2.3 Waveform processing

To test whether waveform differences are able to influence entrainment, we ran the unfiltered spike-triggered average waveform for both proto and arky neurons through typical preprocessing methods. We first obtain an unfiltered, raw spike-triggered average (STA) waveform of the arkys and protos (see fig. 5.1 A & B). Then, a lowpass filter of 500Hz is applied to the waveform, attempting to remove higher frequencies such as that of the waveform (see fig. 5.1 C & D). Next, we downsample the lowpass filtered waveform from 31250 to 496 Hz (see fig. 5.1 E & F). This is a common step when extracting LFP data, as such high-frequency samples are not required to analyze low-frequency oscillations. Next, we bandpass the waveform into the gamma frequency (45-60Hz) (see fig. 5.1 G & H). Finally, we obtain the phase of the spike of the gamma bandpassed waveform by using the Hilbert transform (5.1 I & J). This process is then repeated over all the STAs in arkys and protos to obtain fig. 5.4. For the principal component analysis (PCA), we normalized all values of the waveform to a mean of 0 and a standard deviation of 1. Following this, PCA was performed on the waveforms. We first computed the histogram of coefficients for both arkys and protos for the first, second and third principal components (see fig. 5.3). Next, we projected the waveforms from all the arkys and protos to those represented by their principal components (see second row of fig. 5.3). We then did the same to arkys and protos separately (see third row fig. 5.3).

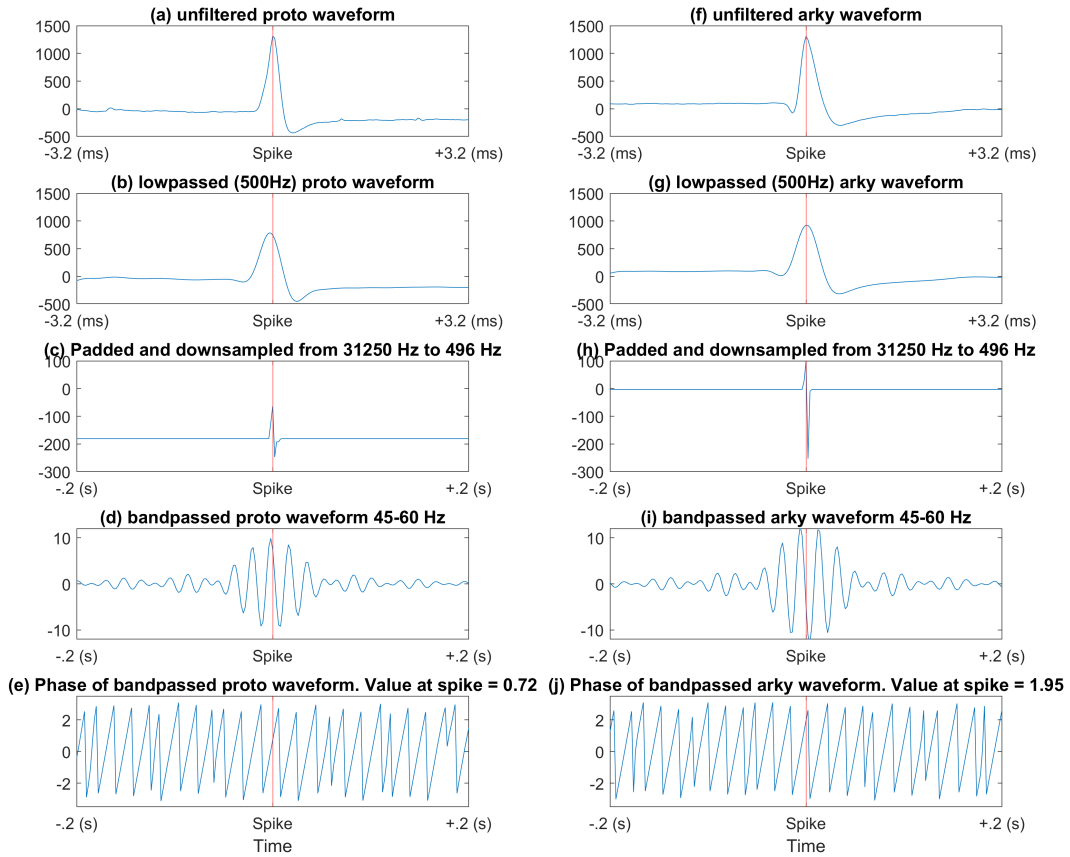


Figure 5.1: Differences in waveform can lead to differences in entrainment. Pre-processing steps for spike-phase entrainment applied to the raw unfiltered waveform of arkys and protos.

5.3 Results

5.3.1 Waveform differences between arkypallidal and prototypical neurons

Figure 5.2 shows the average waveform of all identified putative arky and putative proto neurons 1ms before and 1ms after the spike. We observe that many more arky neurons have lower normalized voltage values both post-spike and pre-spike.

We find that the second principal component best captures the differences between arky and proto cell waveforms, with arkys having a distribution centred around .1 and protos having a distribution centred around -.05 (see fig. 5.2 top row, middle panel). We find significant differences in the distribution between arky and proto coefficients in the second principal component ($p < .0001$, two-sample KS test). The distribution between arky and proto coefficients in the first principal component is significant ($p < .0001$, two-sample KS test). We find that the distribution between arkys and protos coefficients in the third principal component is not significant ($p = .06$, two-sample KS test).

The first principal component consists mainly of the spike, while the second principal component consists of the pre-spike and post-spike dip (compare fig. 5.2 second row, left and middle panels). As the coefficient differences between arkys and protos are significantly different on the second principal component, this explains the differences we observed in the waveforms in fig. 5.2. Again, we see a clear difference between arkys and protos in the second principal component suggesting that arky and proto neurons have different waveforms when recorded with a tetrode.

5.3.2 Differences in waveforms result in different entrainment phases.

As described in the methods, we analyzed spike-phase entrainment for arkys and protos in gamma oscillations. We found 49 arky neurons (44%) that are significantly entrained to 45-55 Hz gamma oscillations (Rayleigh test) and 79 proto neurons (53%) that are significantly entrained to 45-55 Hz gamma oscillations (Rayleigh test). We calculated the mean phase of these significantly entrained arkys and protos and found that arkys are generally entrained to the trough of a gamma wave ($-\pi$) and that protos are entrained about mid-point on the descend-




Figure 5.2: Differences in proto and arky waveforms.. Average waveform of each arky and proto cell (normalised voltage (a.u.)). Arky STA waveforms are on the left, while proto STA waveforms are on the right. Image is sorted by mean values 1ms before the spike.

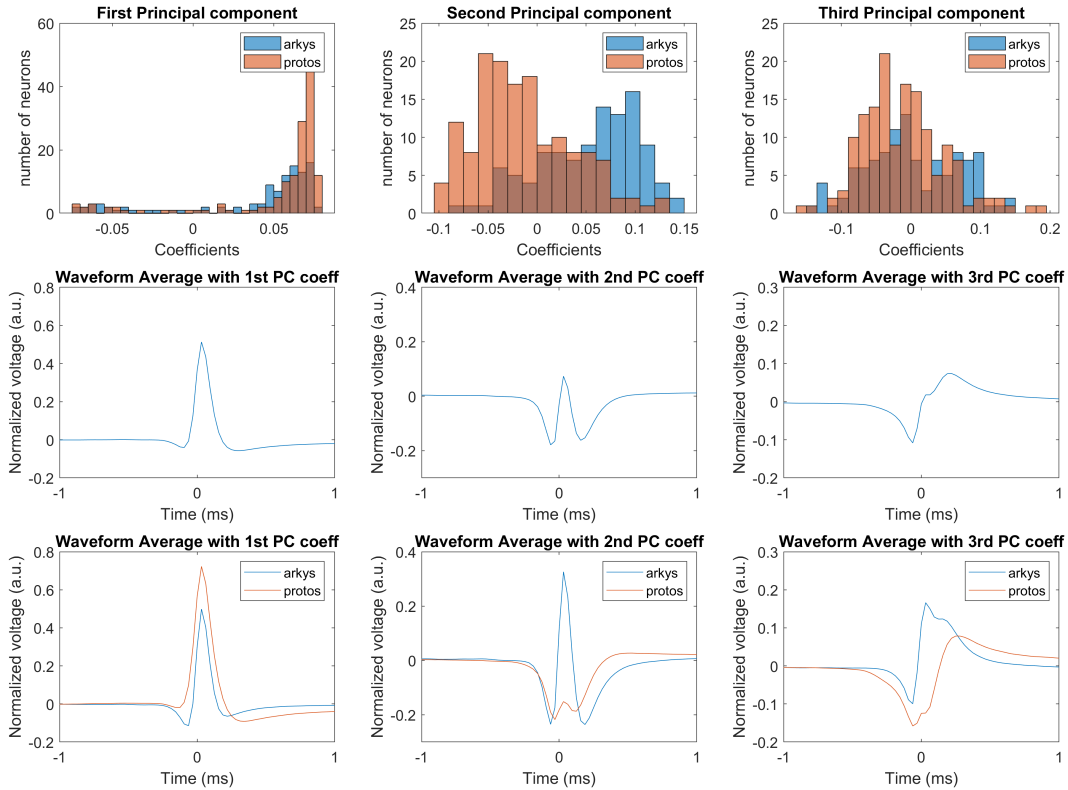


Figure 5.3: Significant differences in the distribution of coefficients for arkys and protos. PCA analysis on raw average waveforms of both arkys and protos. In the first row, the first and second principal components seem to capture some differences between the arkys and protos. In the second and last row, we projected the waveforms from the original data to those represented by the principal components. Arky and proto waveform averages are most distinct for the second principal component (bottom row, middle), with arkys having a higher peak voltage.

ing part of a gamma wave ($-\pi/2$) (see fig. 5.4).

To test whether waveform differences are able to influence entrainment, we processed the unfiltered STA waveform for both proto and arky neurons, similar to how we processed the LFP signal (see methods). The phase of these processed arky and proto waveforms will be where their spike is. We found that we were able to obtain similar distributions as the spike-LFP entrainment above with just the STA waveform (see fig. 5.4). Our results suggest waveform differences in cell types could result in different spike-phase entrainment.

5.4 Discussion

We found differences in waveforms in putative arky and proto neurons. Even though these differences were not found when using other instruments or in other studies (see supplementary figures in Mallet et al. (2016)), it could still be a useful way to identify these neurons when using tetrodes. One reason we observe the waveform differences could be that the tetrodes we use are much smaller than glass electrodes or optrodes. Future studies should work to study waveforms between different types of instruments in GP neurons. However, caution is to be noted as our proto and arky neurons are classified via measures such as slow wave sleep (SWS) firing rate, coefficient of variation (CV) during SWS and SWS rate change. Furthermore, it is difficult to tell whether this waveform distinction is due to anatomical or functional differences or simply due to the recording material or shape.

Mallet et al. (2016) found that 59% of arkys and 26% of protos had a significant amount of entrainment to beta oscillations in freely moving animals given a binomial test. Protos had no significant phase preference, while arkys preferred to fire right before the trough of the beta oscillations. They found few arky and proto neurons significantly entrained to the similar tetrode. This is in contrast to the entrainment to gamma oscillations we found here, using the same dataset. This could be because beta oscillations are of lower frequencies and thus have less contamination of waveforms. Spike waveforms are able to contaminate LFPs to 10Hz theoretically (Waldert et al., 2013), although this has not been tested in the context of spike-entrainment.

Our study highlights the need to use methods such as mean spike removal, in-

terpolation, or bayesian methods (Ray, 2015; Zanos et al., 2011) to process LFPs for studying spike-phase entrainment. These methods use different ways to 'cut' the transient waveform artifact in the LFP data surrounding a short time window around the spike. However, spike-transients from non-detectable spikes can still be present in the LFP (Schomburg et al., 2012). This problem highlights the need to further quantify and study the influence of these spike waveforms on processed LFP data. Finally, while it is also common to use different electrodes for spikes and LFPs, correlations between spikes in a nearby region can still contaminate the LFP and oscillations at the gamma frequency can vary more than $400 \mu m$ apart (Ray, 2015). This suggests that it is difficult to avoid the problem of spike waveform transients in LFP data.

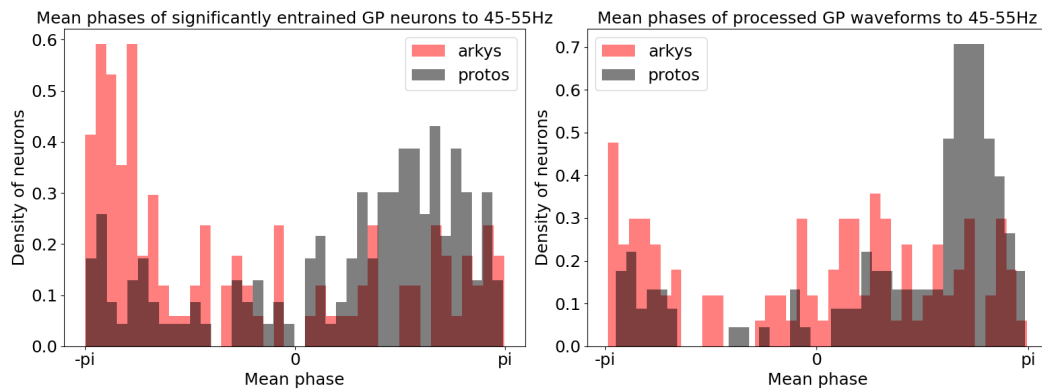


Figure 5.4: Mean phases of significantly entrained GP neurons have similar distributions to phases of processed GP waveforms. Histograms of significantly entrained arky and proto mean phases were obtained with LFP and spike data (left), and phases of arky and proto were obtained with STA (right). Arkys and protos have different distributions of mean phase entrainment to gamma oscillations.

Chapter 6

General discussion

6.1 Overview

In this thesis, we first investigated how to quantify and characterise firing rate ramps found in neural recordings. Specifically, we aimed to investigate these ramps in the basal ganglia and assess regional differences. We then quantified whether single-trial firing rates were more ramp-like or more step-like. Finally, we looked at how different waveforms can contaminate entrainment results.

In chapter 3, we utilised algorithms for ramp detection and changepoint algorithms, some of which have previously been used in wind ramp detection. In order to assess how successful they were at detecting ramps in neural data, we generated simulated datasets where we could control the statistics of the underlying spike trains generating the neural signal. We found that using changepoint algorithms such as PELT to determine ramp start and end points is more effective than other methods to generate piecewise components. Finally, we quantified the sensitivity of hyperparameters in both PELT and `l1tf` and found that small changes in parameters for `l1tf` can result in large changes in results.

In chapter 4, we found ramps in the hold period of a stop signal task in the basal ganglia. The SNr, in particular, had strong negative ramps that extended past the hold period. On the other hand, most regions had positive ramps that started and ended at the go cue. Single-trial units were mostly inconclusive, with many stepping neurons in general. We also developed a method to identify between

stepping and ramping in single trials. We found no correlations between single trials and average ramp characteristics, suggesting that trial-averaged ramps are not represented in the single-trials.

In chapter 5, we investigated the problem of spike bleed-through when studying entrainment. We used putative arky pallidal and prototypical neurons and found waveform differences in them. Furthermore, we demonstrate that these different waveforms can influence spike-phase entrainment.

One interesting observation we found was that the SNr had strong negative ramps that start around the Nose in event after a sharp increase in activity. Similar negative ramps were also previously seen in primate GPi (similar to SNr in rats) (Thura and Cisek, 2017). They suggest that the GPi output offers a time-varying signal that regulates the speed-accuracy trade-off and reflects the increasing need to make a decision and does not contribute to the deliberation process (Thura and Cisek, 2017). The negative SNr ramps we find in the 'hold' period can reflect an increasing expectation of the Go cue. Overall, the SNr seems to be poised in a great position for inhibitory control as it is the main output of the basal ganglia, and the firing activity of many other regions, including the GP, Str, and STN, has projections to the SNr. For example, steps from various regions arriving at different times in the SNr could generate a ramp. The STN also has strong trial-averaged ramps that begin at Nose in. Both regions previously suggested to be involved in action cancellation (STN and SNr) (Schmidt and Berke, 2017) show some clear trial-averaged ramps around the hold period. This suggests that the STN and SNr could also be involved in action restraint. Our results broadly support the idea that STN and SNr could be generally involved in inhibitory control.

Catanese and Jaeger (2021) suggest that the basal ganglia controls up ramping in the thalamus. This is supported by our results where we found down ramping activity in the SNr because the SNr has inhibitory connections to the thalamus. To further support the idea of the basal ganglia controlling thalamic ramps, Tecuapetla et al. (2016) found that inhibition of SNr delays movement time. Catanese and Jaeger (2021) suggested that these ramping activities could be a result of dopamine neurons that project to the striatum. However, the origin of these dopamine ramps is also not yet clear to researchers (see Lerner et al. (2021)).

Dopamine ramps are ramps where dopamine activity increases when an animal approaches a reward. It has been suggested as a way to maintain motivation while working for distal rewards and associating credit to the action that led to a positive outcome (Hamid et al., 2016). These ramps have been found in the ventral tegmental area and the ventral striatum and the dorsomedial striatum (Lerner et al., 2021). The role of these dopamine ramps is controversial and it is still unclear what their function is. One idea is that dopamine ramps can be used as reward prediction errors (Morita and Kato, 2022; Gershman, 2014). The ramps can be the result of backpropagation of RPEs while learning. Another idea is that these dopamine ramps are best understood as motivational cues or as a system for organising goal-directed action planning (Song and Lee, 2020; Sarno et al., 2022). Currently, progress in dopamine ramps research has been hampered by the multitude of recording techniques and the absence of a universally accepted definition of ramping. Our work here on characterising and detecting ramps could also be applied to dopamine ramps in future research. Furthermore, we need to comprehend how different subsets of neurons' dopamine impulses are produced by upstream circuits and translated by downstream circuits.

Many single-trial models, including ours, compare between stepping and ramping (with some including uniform firing rates (Zhao and Kording, 2018)). However, a visual inspection of raster plots of these single trials shows that there are complex combinations of firing rates that could still sum towards a trial-averaged ramp. For example, some ramping neurons had multiple step ups and downs (or bursts) of firing rate within the detected ramp. These complex patterns cannot be easily teased apart with two simple models and require more study or methods to characterise in the context of a trial-averaged ramp. In any case, observing ramps in trial-averaged activity while relating to behaviour can be tricky without understanding how these single trials operate.

Stepping neurons are more common, thus not showing any evidence accumulation. Our results are consistent with that of Zoltowski et al. (2019) where many trial-averaged ramps are stepping in their trials across different datasets. While our task does not have sensory evidence accumulation, ramps could be used as temporal expectation of the Go cue. It is unclear why there would be ramping activity in the trial-averaged activity given stepping neurons. One plausible explanation for this would be that these stepping neurons are representing a state

where the rat is ready to move. The longer the hold period, the more likely that the rat is in a state to move. This could be tested by looking at stepping activity during failed Go trials.

Comparison percentage of neurons across studies is a difficult thing to do. Many studies have a very loose definition of a ramp (e.g. correlation values to a linearly increasing line). While not published in this thesis, we attempted to build a ramp detection algorithm that could detect both linear and non-linear ramps. Characterising the extent to which these non-linear ramps exist and perhaps separating them while analysing behaviour could be useful. These non-linear ramps have been suggested to be found in various other regions (see Zoltowski et al. (2019)).

Many extensions to this work can be done. One simple extension would be to analyse ramps in protos and arkys as well as tasks with various forms of inhibitory control. Furthermore, an attempt to review ramps in past data by comparing their characteristics between studies could also be certainly useful.

Our work builds a foundation for future research to start thinking about how ramps in neuroscience are defined. In the wind ramp literature, different regions have different characterisations of wind ramps, allowing windmill operators to adjust their detection algorithm for optimal energy capture. Similarly, we hope that neuroscience studies could start to better quantify ramps found in various brain regions, for example, the length, startpoint, endpoint, slope, amplitude and signal-to-noise ratio. This will aid us in comparing and understanding ramps better between studies and tasks to inform future research on how to better study or build tasks to observe ramping firing rates.

Bibliography

- Abdi, A., Mallet, N., Mohamed, F. Y., Sharott, A., Dodson, P. D., Nakamura, K. C., Suri, S., Avery, S. V., Larvin, J. T., Garas, F. N., and others (2015). Prototypic and arkypallidal neurons in the dopamine-intact external globus pallidus. *Journal of Neuroscience*, 35(17):6667–6688.
- Alexander, G. E. and Crutcher, M. D. (1990). Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends in neurosciences*, 13(7):266–271.
- Alhourani, A., McDowell, M. M., Randazzo, M. J., Wozny, T. A., Kondylis, E. D., Lipski, W. J., Beck, S., Karp, J. F., Ghuman, A. S., and Richardson, R. M. (2015). Network effects of deep brain stimulation. *Journal of neurophysiology*, 114(4):2105–2117.
- Aminikhanghahi, S. and Cook, D. J. (2017). A survey of methods for time series change point detection. *Knowledge and information systems*, 51(2):339–367.
- Anastassiou, C. A., Perin, R., Markram, H., and Koch, C. (2011). Ephaptic coupling of cortical neurons. *Nature neuroscience*, 14(2):217–223.
- Aristieta, A. and Gittis, A. (2021). Distinct globus pallidus circuits regulate motor and cognitive functions. *Trends in neurosciences*, 44(8):597–599.
- Bari, A. and Robbins, T. W. (2013). Inhibition and impulsivity: behavioral and neural basis of response control. *Progress in neurobiology*, 108:44–79.
- Bartoli, E., Bosking, W., and Foster, B. L. (2020). Seeing visual gamma oscillations in a new light. *Trends in cognitive sciences*, 24(7):501–503.

- Berke, J. D. (2009). Fast oscillations in cortical-striatal networks switch frequency following rewarding events and stimulant drugs. *European Journal of Neuroscience*, 30(5):848–859.
- Beste, C. and Saft, C. (2015). Action selection in a possible model of striatal medium spiny neuron dysfunction: behavioral and EEG data in a patient with benign hereditary chorea. *Brain Structure and Function*, 220(1):221–228.
- Bogacz, R. (2015). Optimal decision making in the cortico-basal-ganglia circuit. In *An Introduction to Model-Based Cognitive Neuroscience*, pages 291–302. Springer.
- Boutros, N. N., Lajiness-O'Neill, R., Zillgitt, A., Richard, A. E., and Bowyer, S. M. (2015). Eeg changes associated with autistic spectrum disorders. *Neuropsychiatric Electrophysiology*, 1(1):1–20.
- Brown, P., Oliviero, A., Mazzone, P., Insola, A., Tonali, P., and Di Lazzaro, V. (2001). Dopamine dependency of oscillations between subthalamic nucleus and pallidum in Parkinson's disease. *Journal of Neuroscience*, 21(3):1033–1038.
- Busch, N. A., Dubois, J., and VanRullen, R. (2009). The phase of ongoing eeg oscillations predicts visual perception. *Journal of neuroscience*, 29(24):7869–7876.
- Buzsáki, G., Anastassiou, C. A., and Koch, C. (2012). The origin of extracellular fields and currents—eeg, ecog, lfp and spikes. *Nature reviews neuroscience*, 13(6):407–420.
- Buzsáki, G. and Wang, X.-J. (2012). Mechanisms of gamma oscillations. *Annual review of neuroscience*, 35:203.
- Calabresi, P., Picconi, B., Tozzi, A., Ghiglieri, V., and Di Filippo, M. (2014). Direct and indirect pathways of basal ganglia: a critical reappraisal. *Nature neuroscience*, 17(8):1022–1030.
- Catanese, J. and Jaeger, D. (2021). Premotor ramping of thalamic neuronal activity is modulated by nigral inputs and contributes to control the timing of action release. *Journal of Neuroscience*, 41(9):1878–1891.
- Chambers, C. D., Garavan, H., and Bellgrove, M. A. (2009). Insights into the neural basis of response inhibition from cognitive and clinical neuroscience. *Neuroscience & biobehavioral reviews*, 33(5):631–646.

- Chan, C. S., Glajch, K. E., Gertler, T. S., Guzman, J. N., Mercer, J. N., Lewis, A. S., Goldberg, A. B., Tkatch, T., Shigemoto, R., Fleming, S. M., and others (2011). HCN channelopathy in external globus pallidus neurons in models of Parkinson's disease. *Nature neuroscience*, 14(1):85–92.
- Cheneka, B. R., Watson, S. J., and Basu, S. (2020). A simple methodology to detect and quantify wind power ramps. *Wind Energy Science*, 5(4):1731–1741.
- Chua, A. D. X. (2017). Significance of globus pallidus external organization on stop signal propagation through the basal ganglia: A computational model. Master's thesis, University of Sheffield.
- Chuhma, N., Tanaka, K. F., Hen, R., and Rayport, S. (2011). Functional connectome of the striatal medium spiny neuron. *Journal of Neuroscience*, 31(4):1183–1192.
- Cole, S. and Voytek, B. (2018). Hippocampal theta bursting and waveform shape reflect ca1 spiking patterns. *BioRxiv*, page 452987.
- Crews, F. T. and Boettiger, C. A. (2009). Impulsivity, frontal lobes and risk for addiction. *Pharmacology Biochemistry and Behavior*, 93(3):237–247.
- Cui, Y., He, Y., Xiong, X., Chen, Z., Li, F., Xu, T., and Zhang, F. (2021). Algorithm for identifying wind power ramp events via novel improved dynamic swinging door. *Renewable Energy*, 171:542–556.
- Diamond, A. (2013). Executive functions. *Annual review of psychology*, 64:135.
- Diesburg, D. A. and Wessel, J. R. (2021). The pause-then-cancel model of human action-stopping: theoretical considerations and empirical evidence. *Neuroscience & Biobehavioral Reviews*, 129:17–34.
- Dodson, P. D., Larvin, J. T., Duffell, J. M., Garas, F. N., Doig, N. M., Kessar, N., Duguid, I. C., Bogacz, R., Butt, S. J., and Magill, P. J. (2015). Distinct developmental origins manifest in the specialized encoding of movement by adult neurons of the external globus pallidus. *Neuron*, 86(2):501–513.
- Dong, J., Hawes, S., Wu, J., Le, W., and Cai, H. (2021). Connectivity and functionality of the globus pallidus externa under normal conditions and parkinson's disease. *Front. Neural Circuits*, 15:645287.

- Donnelly, N. A., Paulsen, O., Robbins, T. W., and Dalley, J. W. (2015). Ramping single unit activity in the medial prefrontal cortex and ventral striatum reflects the onset of waiting but not imminent impulsive actions. *European Journal of Neuroscience*, 41(12):1524–1537.
- Eagle, D. M., Wong, J. C., Allan, M. E., Mar, A. C., Theobald, D. E., and Robbins, T. W. (2011). Contrasting roles for dopamine d1 and d2 receptor subtypes in the dorsomedial striatum but not the nucleus accumbens core during behavioral inhibition in the stop-signal task in rats. *Journal of Neuroscience*, 31(20):7349–7356.
- Emmons, E. B., De Corte, B. J., Kim, Y., Parker, K. L., Matell, M. S., and Narayanan, N. S. (2017). Rodent medial frontal control of temporal processing in the dorso-medial striatum. *Journal of Neuroscience*, 37(36):8718–8733.
- Engel, A. K. and Fries, P. (2010). Beta-band oscillations—signalling the status quo? *Current opinion in neurobiology*, 20(2):156–165.
- Fee, M. S., Mitra, P. P., and Kleinfeld, D. (1996). Variability of extracellular spike waveforms of cortical neurons. *Journal of neurophysiology*, 76(6):3823–3833.
- Fife, K. H., Gutierrez-Reed, N. A., Zell, V., Bailly, J., Lewis, C. M., Aron, A. R., and Hnasko, T. S. (2017). Causal role for the subthalamic nucleus in interrupting behavior. *eLife*, 6.
- Florita, A., Hodge, B.-M., and Orwig, K. (2013). Identifying wind and solar ramping events. In *2013 IEEE Green Technologies Conference (GreenTech)*, pages 147–152. IEEE.
- Frank, M. J. (2006). Hold your horses: a dynamic computational role for the subthalamic nucleus in decision making. *Neural networks*, 19(8):1120–1136.
- Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in cognitive sciences*, 9(10):474–480.
- Fries, P. (2015). Rhythms for cognition: communication through coherence. *Neuron*, 88(1):220–235.
- Gallego-Castillo, C., Cuerva-Tejero, A., and Lopez-Garcia, O. (2015). A review on the recent history of wind power ramp forecasting. *Renewable and Sustainable Energy Reviews*, 52:1148–1157.

- Gerfen, C. R. and Surmeier, D. J. (2011). Modulation of striatal projection systems by dopamine. *Annual review of neuroscience*, 34:441.
- Gershman, S. J. (2014). Dopamine ramps are a consequence of reward prediction errors. *Neural computation*, 26(3):467–471.
- Glajch, K. E., Kolver, D. A., Hegeman, D. J., Cui, Q., Xenias, H. S., Augustine, E. C., Hernández, V. M., Verma, N., Huang, T. Y., Luo, M., and others (2016). Npas1+ pallidal neurons target striatal projection neurons. *Journal of Neuroscience*, 36(20):5472–5488.
- Goenner, L., Maith, O., Koulouri, I., Baladron, J., and Hamker, F. H. (2021). A spiking model of basal ganglia dynamics in stopping behavior supported by arkypallidal neurons. *European Journal of Neuroscience*, 53(7):2296–2321.
- Goldberg, J. H., Farries, M. A., and Fee, M. S. (2013). Basal ganglia output to the thalamus: still a paradox. *Trends in neurosciences*, 36(12):695–705.
- Goldberg, J. H. and Fee, M. S. (2012). A cortical motor nucleus drives the basal ganglia-recipient thalamus in singing birds. *Nature neuroscience*, 15(4):620–627.
- Gonzalez, V., Cif, L., Biolsi, B., Garcia-Ptacek, S., Seychelles, A., Sanrey, E., Descours, I., Coubes, C., de Moura, A.-M. R., Corlobe, A., and others (2014). Deep brain stimulation for Huntington’s disease: long-term results of a prospective open-label study. *Journal of neurosurgery*, 121(1):114–122.
- Gu, B.-M. and Berke, J. D. (2022). Altered basal ganglia output during self-restraint. *bioRxiv*.
- Gurney, K., Prescott, T. J., and Redgrave, P. (2001a). A computational model of action selection in the basal ganglia. i. a new functional anatomy. *Biological cybernetics*, 84(6):401–410.
- Gurney, K., Prescott, T. J., and Redgrave, P. (2001b). A computational model of action selection in the basal ganglia. ii. analysis and simulation of behaviour. *Biological cybernetics*, 84(6):411–423.
- Gut-Fayand, A., Dervaux, A., Olié, J.-P., Lôo, H., Poirier, M.-F., and Krebs, M.-O. (2001). Substance abuse and suicidality in schizophrenia: a common risk factor linked to impulsivity. *Psychiatry research*, 102(1):65–72.

- Hamid, A. A., Pettibone, J. R., Mabrouk, O. S., Hetrick, V. L., Schmidt, R., Vander Weele, C. M., Kennedy, R. T., Aragona, B. J., and Berke, J. D. (2016). Mesolimbic dopamine signals the value of work. *Nature neuroscience*, 19(1):117–126.
- Han, K.-S., Guo, C., Chen, C. H., Witter, L., Osorno, T., and Regehr, W. G. (2018). Ephaptic coupling promotes synchronous firing of cerebellar purkinje cells. *Neuron*, 100(3):564–578.
- Hannesdóttir, Á. and Kelly, M. (2019). Detection and characterization of extreme wind speed ramps. *Wind Energy Science*, 4(3):385–396.
- Hegeman, D. J., Hong, E. S., Hernandez, V. M., and Chan, C. S. (2016). The external globus pallidus: progress and perspectives. *European Journal of Neuroscience*, 43(10):1239–1265.
- Herreras, O. (2016). Local field potentials: myths and misunderstandings. *Frontiers in neural circuits*, 10:101.
- Inase, M., Buford, J. A., and Anderson, M. E. (1996). Changes in the control of arm position, movement, and thalamic discharge during local inactivation in the globus pallidus of the monkey. *Journal of Neurophysiology*, 75(3):1087–1104.
- Jackson, B., Scargle, J. D., Barnes, D., Arabhi, S., Alt, A., Gioumousis, P., Gwin, E., Sangtrakulcharoen, P., Tan, L., and Tsai, T. T. (2005). An algorithm for optimal partitioning of data on an interval. *IEEE Signal Processing Letters*, 12(2):105–108.
- Jensen, O. and Lisman, J. E. (2000). Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase coding. *Journal of neurophysiology*, 83(5):2602–2609.
- Kiani, R., Hanks, T. D., and Shadlen, M. N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *Journal of Neuroscience*, 28(12):3017–3029.
- Kiani, R. and Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *science*, 324(5928):759–764.
- Killick, R., Fearnhead, P., and Eckley, I. A. (2012). Optimal detection of change-points with a linear computational cost. *Journal of the American Statistical Association*, 107(500):1590–1598.

- Kita, H. and Kita, S. T. (1994). The morphology of globus pallidus projection neurons in the rat: an intracellular staining study. *Brain research*, 636(2):308–319.
- Klausberger, T. and Somogyi, P. (2008). Neuronal diversity and temporal dynamics: the unity of hippocampal circuit operations. *Science*, 321(5885):53–57.
- Knyazev, G. G. (2007). Motivation, emotion, and their inhibitory control mirrored in brain oscillations. *Neuroscience & Biobehavioral Reviews*, 31(3):377–395.
- Kravitz, A. V., Freeze, B. S., Parker, P. R., Kay, K., Thwin, M. T., Deisseroth, K., and Kreitzer, A. C. (2010). Regulation of parkinsonian motor behaviors by optogenetic control of basal ganglia circuitry. *Nature*, 466(7306):622.
- Lachenmayer, M. L., Mürset, M., Antih, N., Debove, I., Muellner, J., Bompert, M., Schlaeppli, J.-A., Nowacki, A., You, H., Michelis, J. P., et al. (2021). Subthalamic and pallidal deep brain stimulation for parkinson’s disease—meta-analysis of outcomes. *NPJ Parkinson’s disease*, 7(1):1–10.
- Lanciego, J. L., Luquin, N., and Obeso, J. A. (2012). Functional neuroanatomy of the basal ganglia. *Cold Spring Harbor perspectives in medicine*, 2(12):a009621.
- Latimer, K. W., Yates, J. L., Meister, M. L., Huk, A. C., and Pillow, J. W. (2015). Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, 349(6244):184–187.
- Lerner, T. N., Holloway, A. L., and Seiler, J. L. (2021). Dopamine, updated: reward prediction error and beyond. *Current opinion in neurobiology*, 67:123–130.
- Leventhal, D. K., Gage, G. J., Schmidt, R., Pettibone, J. R., Case, A. C., and Berke, J. D. (2012). Basal ganglia beta oscillations accompany cue utilization. *Neuron*, 73(3):523–536.
- Lijffijt, M., Kenemans, J. L., Verbaten, M. N., and van Engeland, H. (2005). *A meta-analytic review of stopping performance in attention-deficit/hyperactivity disorder: deficient inhibitory motor control?* American Psychological Association.
- Lindén, H., Pettersen, K. H., and Einevoll, G. T. (2010). Intrinsic dendritic filtering gives low-pass power spectra of local field potentials. *Journal of computational neuroscience*, 29(3):423–444.

- Lindén, H., Tetzlaff, T., Potjans, T. C., Pettersen, K. H., Grün, S., Diesmann, M., and Einevoll, G. T. (2011). Modeling the spatial reach of the lfp. *Neuron*, 72(5):859–872.
- Lipszyc, J. and Schachar, R. (2010). Inhibitory control and psychopathology: a meta-analysis of studies using the stop signal task. *Journal of the International Neuropsychological Society*, 16(6):1064–1076.
- Logan, G. D. (1994). On the ability to inhibit thought and action: A users' guide to the stop signal paradigm.
- London, T. D., Licholai, J. A., Szczot, I., Ali, M. A., LeBlanc, K. H., Fobbs, W. C., and Kravitz, A. V. (2018). Coordinated ramping of dorsal striatal pathways preceding food approach and consumption. *Journal of Neuroscience*, 38(14):3547–3558.
- Magnusson, J. L. and Leventhal, D. K. (2021). Revisiting the “paradox of stereotaxic surgery”: insights into basal ganglia-thalamic interactions. *Frontiers in Systems Neuroscience*, page 91.
- Mallet, N., Micklem, B. R., Henny, P., Brown, M. T., Williams, C., Bolam, J. P., Nakamura, K. C., and Magill, P. J. (2012). Dichotomous organization of the external globus pallidus. *Neuron*, 74(6):1075–1086.
- Mallet, N., Pogosyan, A., Márton, L. F., Bolam, J. P., Brown, P., and Magill, P. J. (2008). Parkinsonian beta oscillations in the external globus pallidus and their relationship with subthalamic nucleus activity. *Journal of neuroscience*, 28(52):14245–14258.
- Mallet, N., Schmidt, R., Leventhal, D., Chen, F., Amer, N., Boraud, T., and Berke, J. D. (2016). Arky pallidal cells send a stop signal to striatum. *Neuron*, 89(2):308–316.
- Mar, K., Townes, P., Pechlivanoglou, P., Arnold, P., and Schachar, R. (2022). Obsessive compulsive disorder and response inhibition: Meta-analysis of the stop-signal task. *Journal of Psychopathology and Clinical Science*, 131(2):152.
- Meister, M. L., Hennig, J. A., and Huk, A. C. (2013). Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. *Journal of Neuroscience*, 33(6):2254–2267.

- Milardi, D., Gaeta, M., Marino, S., Arrigo, A., Vaccarino, G., Mormina, E., Rizzo, G., Milazzo, C., Finocchio, G., Baglieri, A., et al. (2015). Basal ganglia network by constrained spherical deconvolution: A possible cortico-pallidal pathway? *Movement Disorders*, 30(3):342–349.
- Mirzaei, A., Kumar, A., Leventhal, D., Mallet, N., Aertsen, A., Berke, J., and Schmidt, R. (2017). Sensorimotor processing in the basal ganglia leads to transient beta oscillations during behavior. *Journal of Neuroscience*, 37(46):11220–11232.
- Morita, K. and Kato, A. (2022). Dopamine ramps for accurate value learning under uncertainty. *Trends in Neurosciences*.
- Nagel, S. J., Machado, A. G., Gale, J. T., Lobel, D. A., and Pandya, M. (2015). Preserving cortico-striatal function: deep brain stimulation in Huntington’s disease. *Frontiers in systems neuroscience*, 9.
- Nambu, A., Chiken, S., Shashidharan, P., Nishibayashi, H., Ogura, M., Kakishita, K., Tanaka, S., Tachibana, Y., Kita, H., and Itakura, T. (2011). Reduced pallidal output causes dystonia. *Frontiers in systems neuroscience*, 5.
- Nambu, A. and Llinas, R. (1997). Morphology of globus pallidus neurons: its correlation with electrophysiology in guinea pig brain slices. *Journal of Comparative Neurology*, 377(1):85–94.
- Nambu, A., Tokuno, H., and Takada, M. (2002). Functional significance of the cortico–subthalamo–pallidal ‘hyperdirect’ pathway. *Neuroscience research*, 43(2):111–117.
- Narayanan, N. S. (2016). Ramping activity is a cortical mechanism of temporal control of action. *Current opinion in behavioral sciences*, 8:226–230.
- Nelson, A. B. and Kreitzer, A. C. (2014). Reassessing models of basal ganglia function and dysfunction. *Annual review of neuroscience*, 37:117–135.
- Nevado-Holgado, A. J., Mallet, N., Magill, P. J., and Bogacz, R. (2014). Effective connectivity of the subthalamic nucleus–globus pallidus network during Parkinsonian oscillations. *The Journal of physiology*, 592(7):1429–1455.

- Nigg, J. T. (2001). Is adhd a disinhibitory disorder? *Psychological bulletin*, 127(5):571.
- Pesaran, B., Vinck, M., Einevoll, G. T., Sirota, A., Fries, P., Siegel, M., Truccolo, W., Schroeder, C. E., and Srinivasan, R. (2018). Investigating large-scale brain dynamics using field potential recordings: analysis and interpretation. *Nature neuroscience*, 21(7):903–919.
- Priori, A., Foffani, G., Rossi, L., and Marceglia, S. (2013). Adaptive deep brain stimulation (adbs) controlled by local field potential oscillations. *Experimental neurology*, 245:77–86.
- Ray, S. (2015). Challenges in the quantification and interpretation of spike-lfp relationships. *Current opinion in neurobiology*, 31:111–118.
- Ray, S. (2022). Spike–gamma phase relationship in the visual cortex.
- Ray, S. and Maunsell, J. H. (2015). Do gamma oscillations play a role in cerebral cortex? *Trends in cognitive sciences*, 19(2):78–85.
- Robbins, T. W., Gillan, C. M., Smith, D. G., de Wit, S., and Ersche, K. D. (2012). Neurocognitive endophenotypes of impulsivity and compulsivity: towards dimensional psychiatry. *Trends in cognitive sciences*, 16(1):81–91.
- Roitman, J. D. and Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of neuroscience*, 22(21):9475–9489.
- Rosin, B., Slovik, M., Mitelman, R., Rivlin-Etzion, M., Haber, S. N., Israel, Z., Vaadia, E., and Bergman, H. (2011). Closed-loop deep brain stimulation is superior in ameliorating parkinsonism. *Neuron*, 72(2):370–384.
- Saggar, M., Sporns, O., Gonzalez-Castillo, J., Bandettini, P. A., Carlsson, G., Glover, G., and Reiss, A. L. (2018). Towards a new approach to reveal dynamical organization of the brain using topological data analysis. *Nature communications*, 9(1):1–14.
- Sarno, S., Beirán, M., Falcó-Roget, J., Diaz-deLeon, G., Rossi-Pool, R., Romo, R., and Parga, N. (2022). Dopamine firing plays a dual role in coding reward prediction errors and signaling motivation in a working memory task. *Proceedings of the National Academy of Sciences*, 119(2):e2113311119.

- Schmidt, R. and Berke, J. D. (2017). A Pause-then-Cancel model of stopping: evidence from basal ganglia neurophysiology. *Phil. Trans. R. Soc. B*, 372(1718):20160202.
- Schmidt, R., Leventhal, D. K., Mallet, N., Chen, F., and Berke, J. D. (2013). Canceling actions involves a race between basal ganglia pathways. *Nature neuroscience*, 16(8):1118–1124.
- Schomburg, E. W., Anastassiou, C. A., Buzsáki, G., and Koch, C. (2012). The spiking component of oscillatory extracellular potentials in the rat hippocampus. *Journal of Neuroscience*, 32(34):11798–11811.
- Schwab, B. C., Kase, D., Zimnik, A., Rosenbaum, R., Codianni, M. G., Rubin, J. E., and Turner, R. S. (2020). Neural activity during a simple reaching task in macaques is counter to gating and rebound in basal ganglia–thalamic communication. *PLoS biology*, 18(10):e3000829.
- Sevlian, R. and Rajagopal, R. (2013). Detection and statistics of wind power ramps. *IEEE Transactions on Power Systems*, 28(4):3610–3620.
- Simen, P., Balci, F., deSouza, L., Cohen, J. D., and Holmes, P. (2011). A model of interval timing by neural integration. *Journal of Neuroscience*, 31(25):9238–9253.
- Smith, M. A. and Kohn, A. (2008). Spatial and temporal scales of neuronal correlation in primary visual cortex. *Journal of Neuroscience*, 28(48):12591–12603.
- Song, M. R. and Lee, S. W. (2020). Dynamic resource allocation during reinforcement learning accounts for ramping and phasic dopamine activity. *Neural Networks*, 126:95–107.
- Spaak, E., de Lange, F. P., and Jensen, O. (2014). Local entrainment of alpha oscillations by visual stimuli causes cyclic modulation of perception. *Journal of Neuroscience*, 34(10):3536–3544.
- Stevens, L., Verdejo-García, A., Goudriaan, A. E., Roeyers, H., Dom, G., and Vanderplasschen, W. (2014). Impulsivity as a vulnerability factor for poor addiction treatment outcomes: a review of neurocognitive findings among individuals with substance use disorders. *Journal of substance abuse treatment*, 47(1):58–72.

- Tecuapetla, F., Jin, X., Lima, S. Q., and Costa, R. M. (2016). Complementary contributions of striatal projection pathways to action initiation and execution. *Cell*, 166(3):703–715.
- Thura, D. and Cisek, P. (2017). The basal ganglia do not select reach targets but control the urgency of commitment. *Neuron*, 95(5):1160–1170.
- Truong, C., Oudre, L., and Vayatis, N. (2020). Selective review of offline change point detection methods. *Signal Processing*, 167:107299.
- Uhlhaas, P. J. and Singer, W. (2010). Abnormal neural oscillations and synchrony in schizophrenia. *Nature reviews neuroscience*, 11(2):100–113.
- van Wingerden, M., Vinck, M., Lankelma, J. V., and Pennartz, C. M. (2010). Learning-associated gamma-band phase-locking of action–outcome selective neurons in orbitofrontal cortex. *Journal of Neuroscience*, 30(30):10025–10038.
- Verbruggen, F. and Logan, G. D. (2009). Models of response inhibition in the stop-signal and stop-change paradigms. *Neuroscience & Biobehavioral Reviews*, 33(5):647–661.
- Waldert, S., Lemon, R. N., and Kraskov, A. (2013). Influence of spiking activity on cortical local field potentials. *The Journal of physiology*, 591(21):5291–5303.
- Wei, W. and Wang, X.-J. (2016). Inhibitory Control in the Cortico-Basal Ganglia-Thalamocortical Loop: Complex Regulation and Interplay with Memory and Decision Processes. *Neuron*, 92(5):1093–1105.
- Wilson, C. J., Higgs, M. H., Simmons, D. V., and Morales, J. C. (2018). Oscillations and spike entrainment. *F1000Research*, 7.
- Winstanley, C. A., Eagle, D. M., and Robbins, T. W. (2006). Behavioral models of impulsivity in relation to adhd: translation between clinical and preclinical studies. *Clinical psychology review*, 26(4):379–395.
- Womelsdorf, T., Ardid, S., Everling, S., and Valiante, T. A. (2014). Burst firing synchronizes prefrontal and anterior cingulate cortex during attentional control. *Current Biology*, 24(22):2613–2621.

- Yasukawa, T., Kita, T., Xue, Y., and Kita, H. (2004). Rat intralaminar thalamic nuclei projections to the globus pallidus: a biotinylated dextran amine anterograde tracing study. *Journal of Comparative Neurology*, 471(2):153–167.
- Zanos, T. P., Mineault, P. J., and Pack, C. C. (2011). Removal of spurious correlations between spikes and local field potentials. *Journal of neurophysiology*, 105(1):474–486.
- Zhao, X. and Kording, K. P. (2018). Rate fluctuations not steps dominate lip activity during decision-making. *bioRxiv*, page 249672.
- Zoltowski, D. M., Latimer, K. W., Yates, J. L., Huk, A. C., and Pillow, J. W. (2019). Discrete stepping and nonlinear ramping dynamics underlie spiking responses of lip neurons during decision-making. *Neuron*, 102(6):1249–1258.

Appendices

Appendix A

Chapter 3: Quantification of firing rate ramps

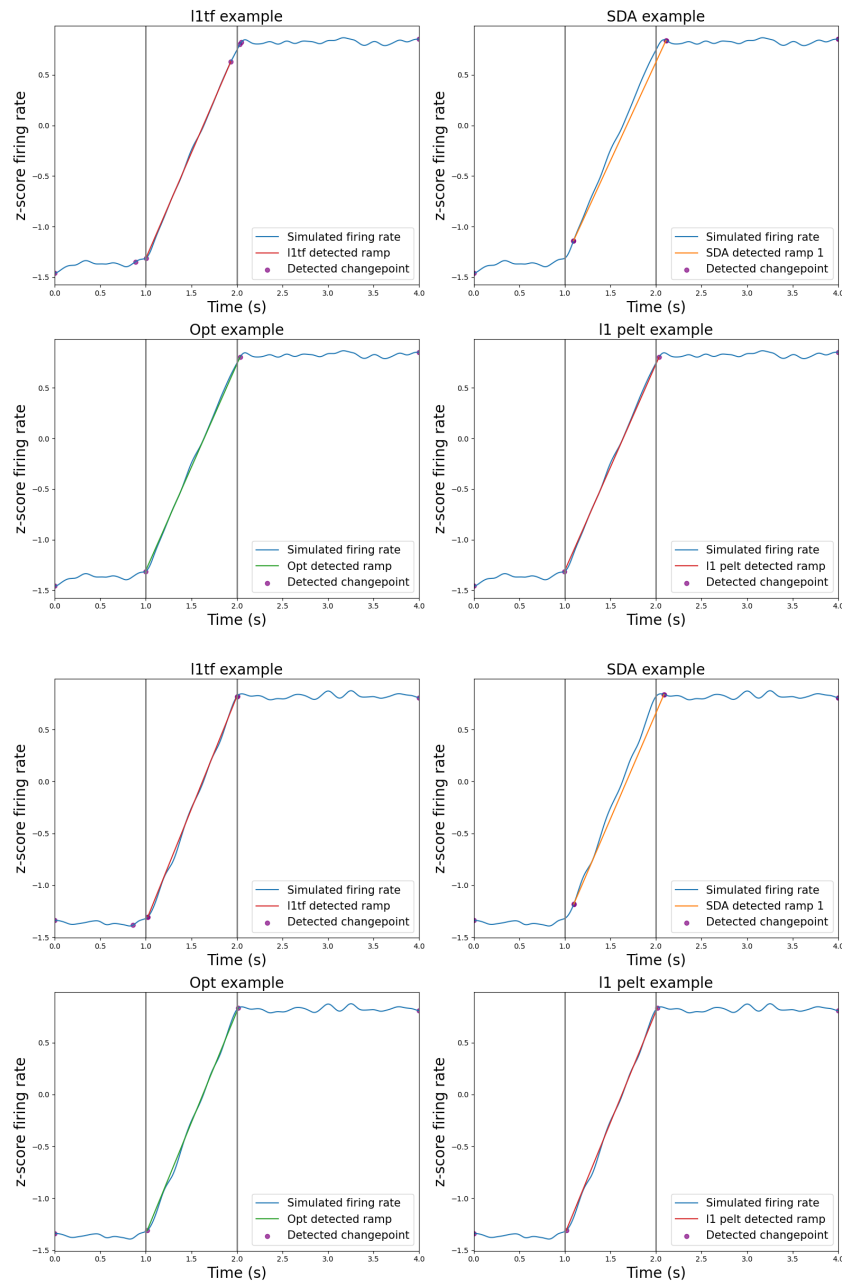


Figure A.1: Example low-noise simulations with l1tf, SDA, Opt and PELT. Gaussian noise with standard deviation of 5ms was used for the simulation. Upper four panels show one simulated ramp and lower four panels show another simulated ramp. Upper left panel shows the detected ramps using l1tf. Upper right panel shows the SDA algorithm applied on the same simulation. Lower left panel shows the Opt changepoint algorithm with the number of breakpoints given. Lower right panel shows the PELT algorithm.

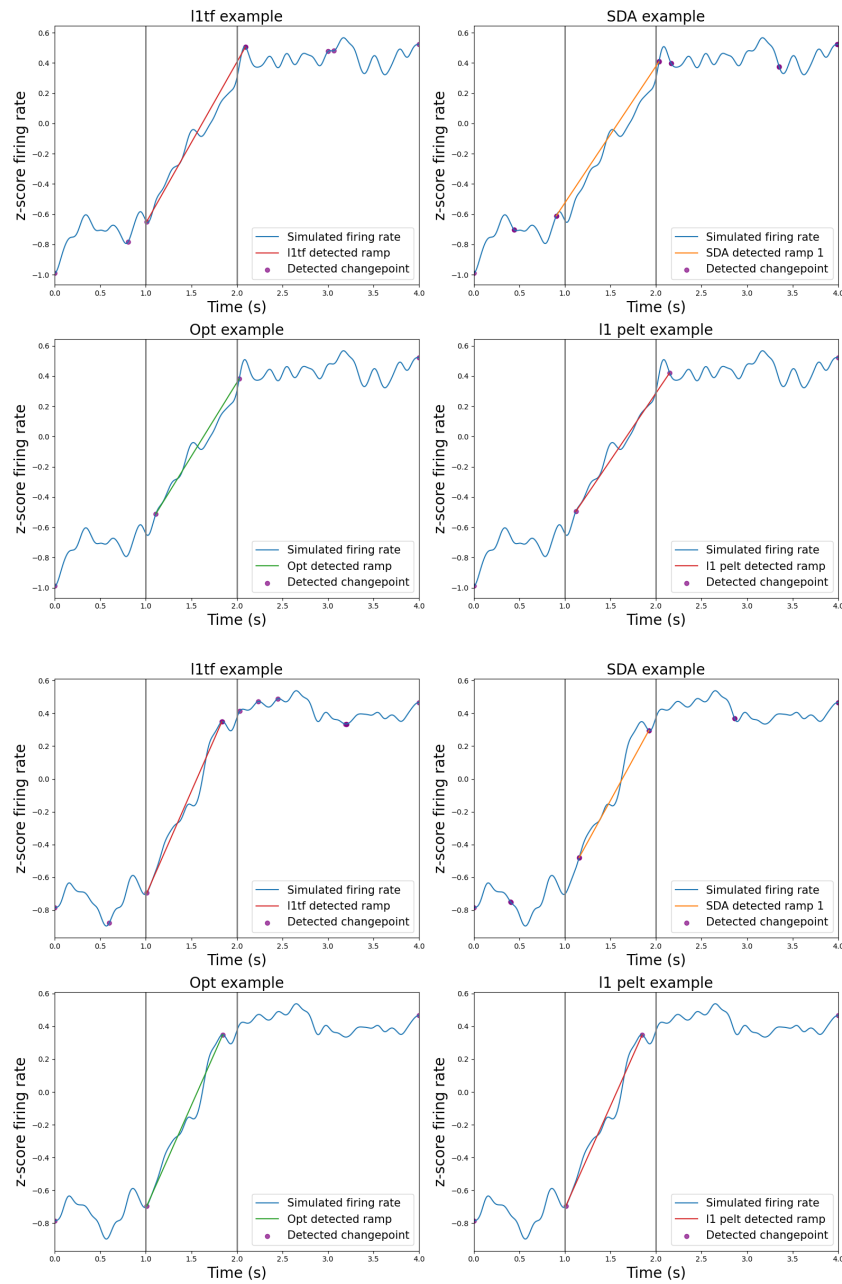


Figure A.2: Example medium-noise simulations with l1tf, SDA, Opt and PELT. Gaussian noise with standard deviation of 30ms was used for the simulation. Upper four panels show one simulated ramp and lower four panels show another simulated ramp. Upper left panel shows the detected ramps using l1tf. Upper right panel shows the SDA algorithm applied on the same simulation. Lower left panel shows the Opt changepoint algorithm with the number of breakpoints given. Lower right panel shows the PELT algorithm.

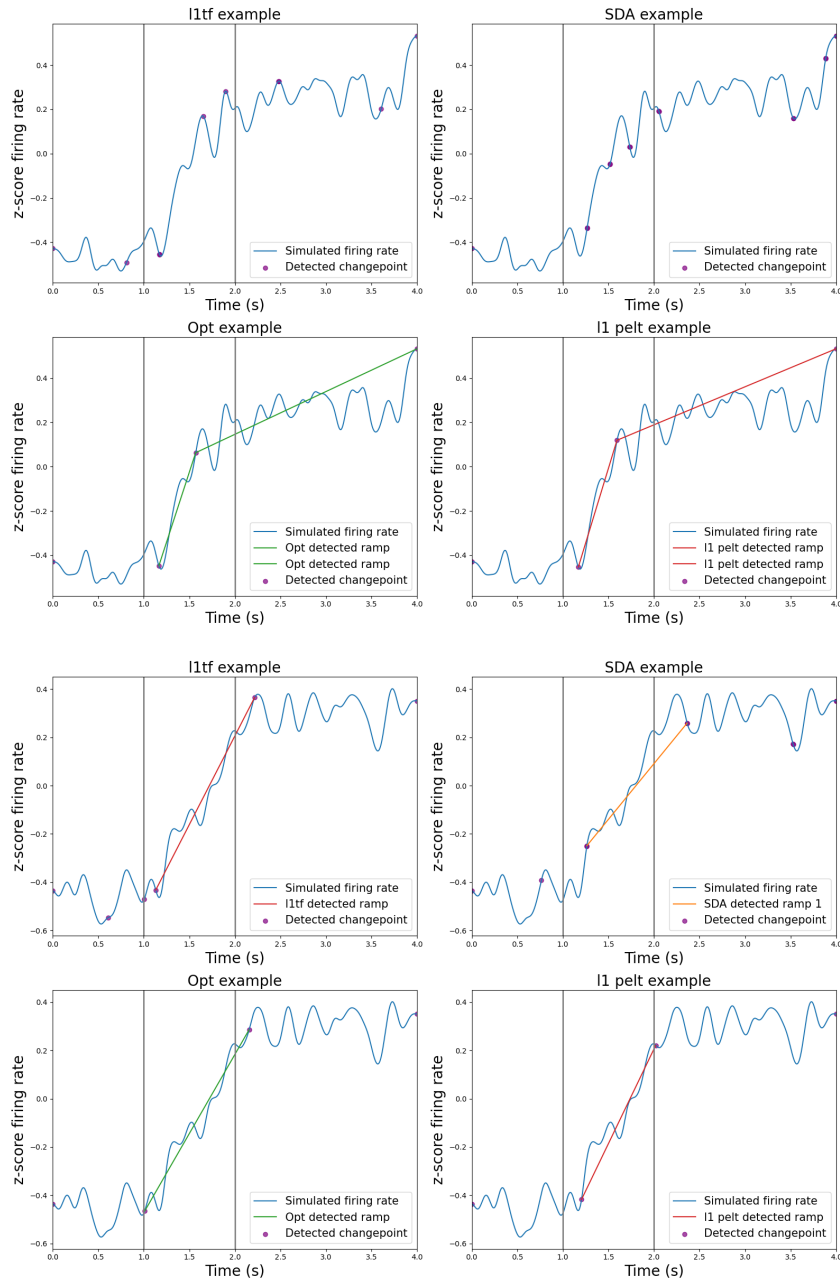


Figure A.3: Example high-noise simulations with l1tf, SDA, Opt and PELT. Gaussian noise with standard deviation of 50ms was used for the simulation. Upper four panels show one simulated ramp and lower four panels show another simulated ramp. Upper left panel shows the detected ramps using l1tf. Upper right panel shows the SDA algorithm applied on the same simulation. Lower left panel shows the Opt changepoint algorithm with the number of breakpoints given. Lower right panel shows the PELT algorithm.

Appendix B

Chapter 4: Properties of ramping activity in the basal ganglia

B.1 Detected ramps aligned to Nose in

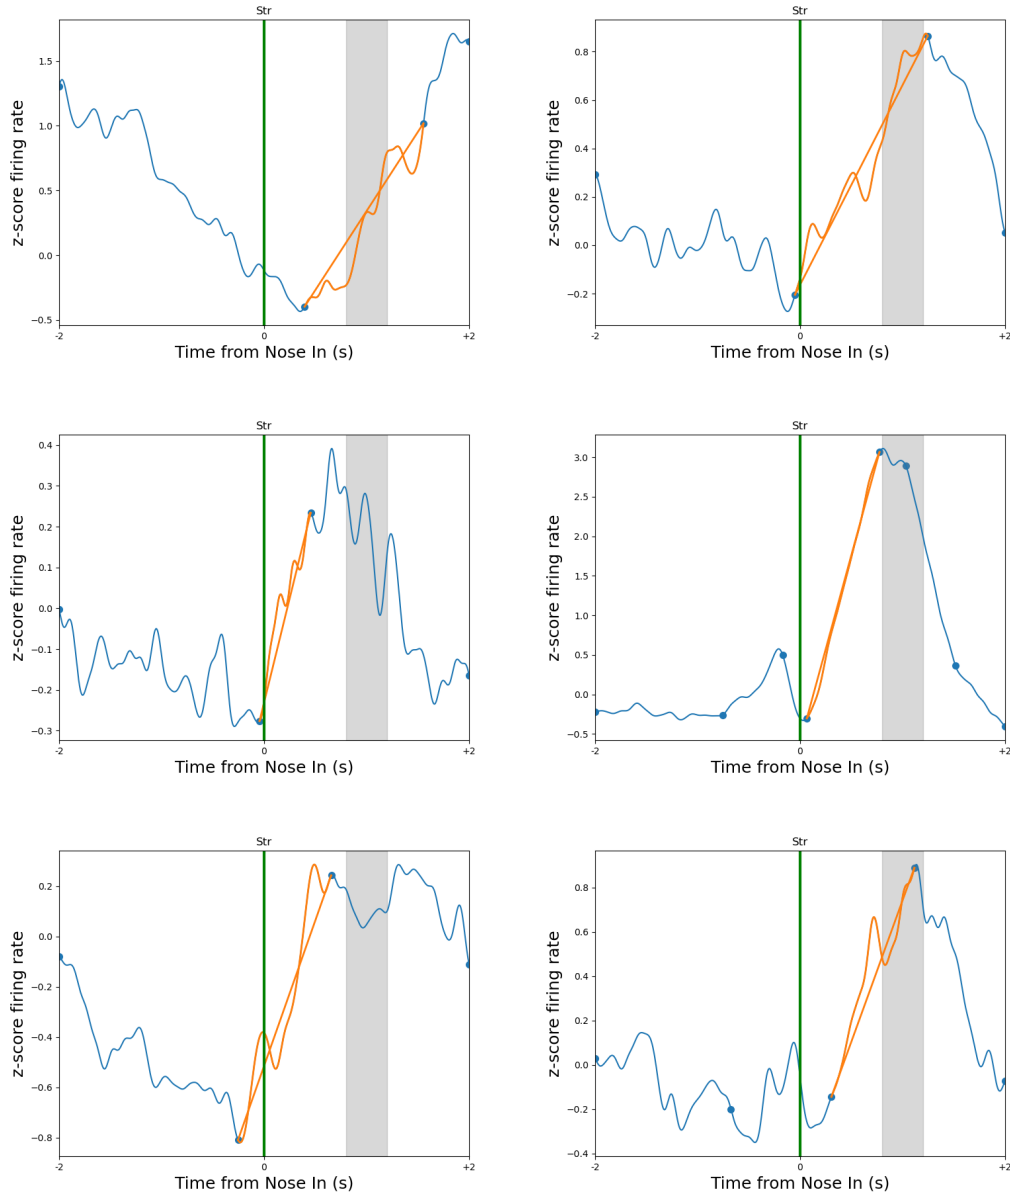


Figure B.1: Examples of ramp detection using the PELT algorithm for the Str. Orange section indicates detected ramp. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

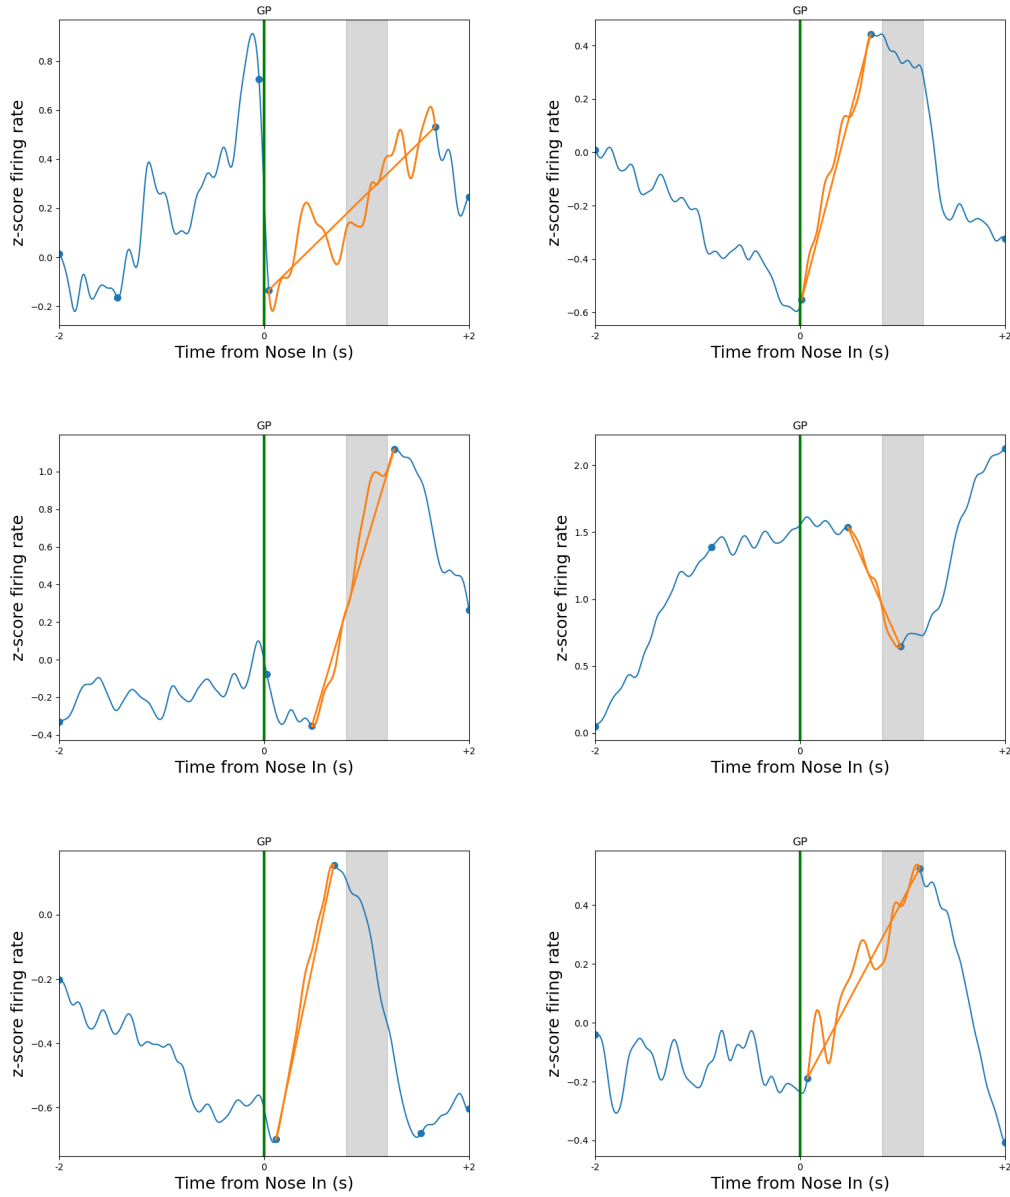


Figure B.2: Examples of ramp detection using the PELT algorithm for the GP. Orange section indicates detected ramp. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

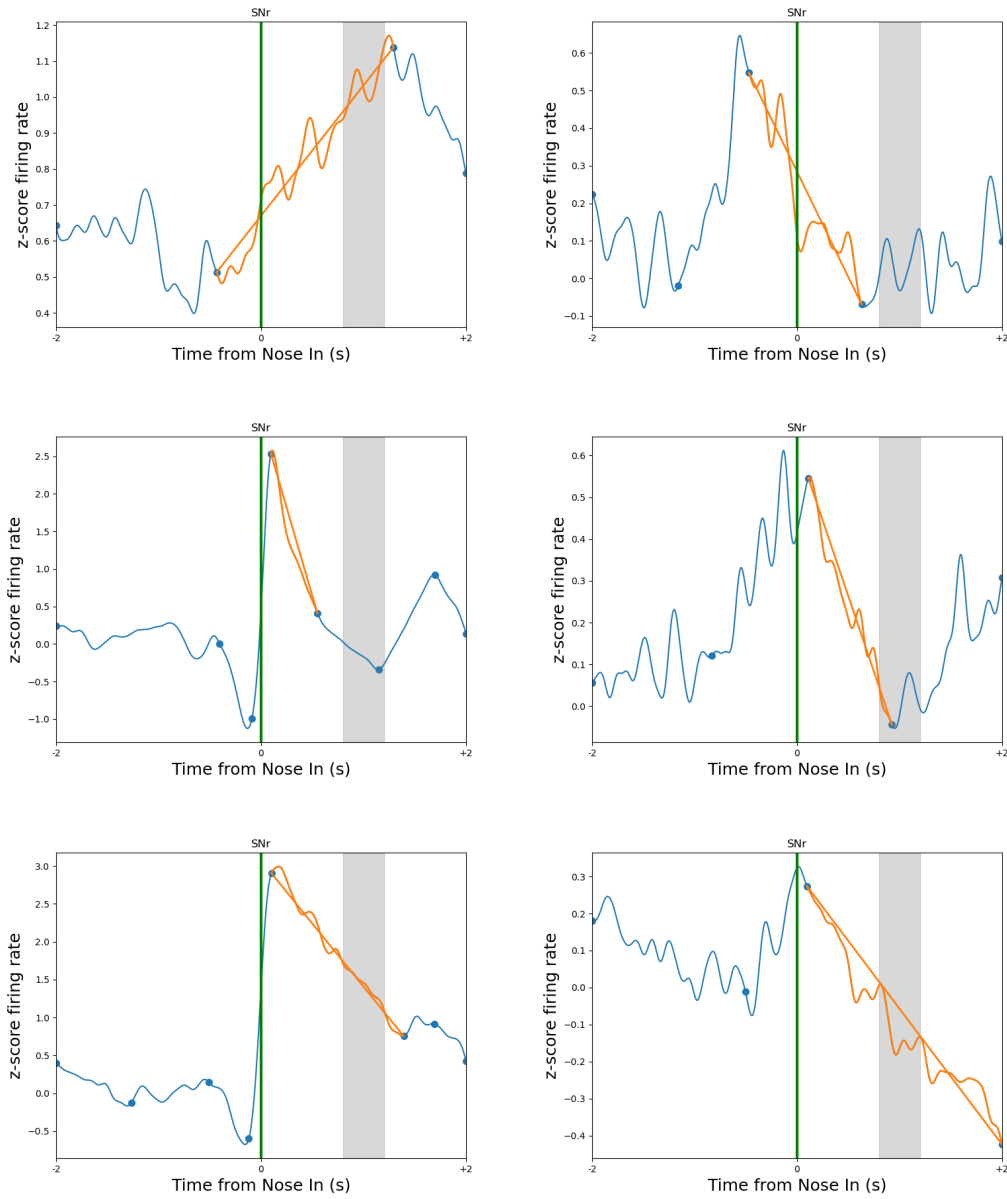


Figure B.3: Examples of ramp detection using the PELT algorithm for the SNr. Orange section indicates detected ramp. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

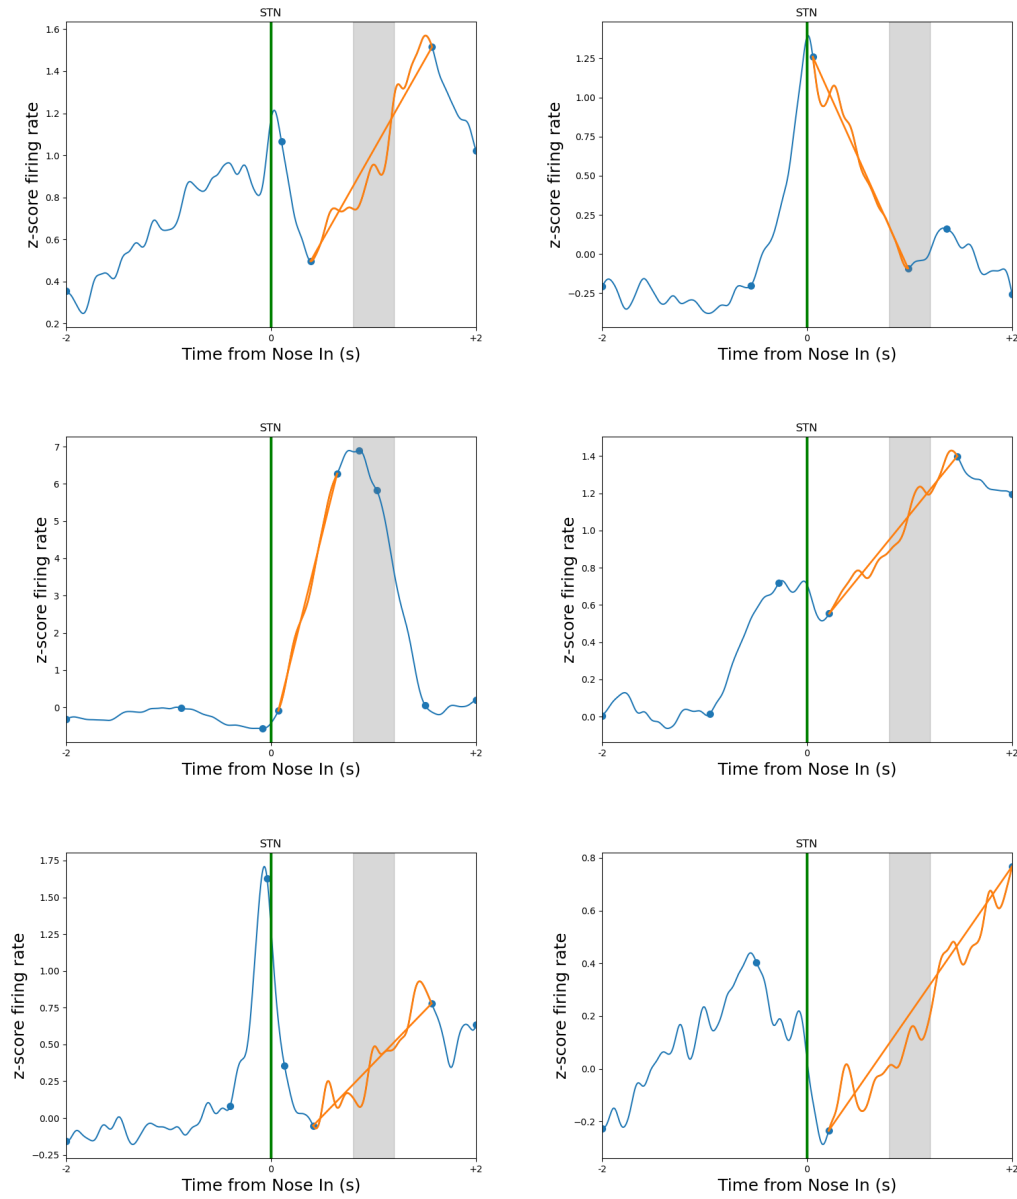


Figure B.4: Examples of ramp detection using the PELT algorithm for the STN. Orange section indicates detected ramp. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

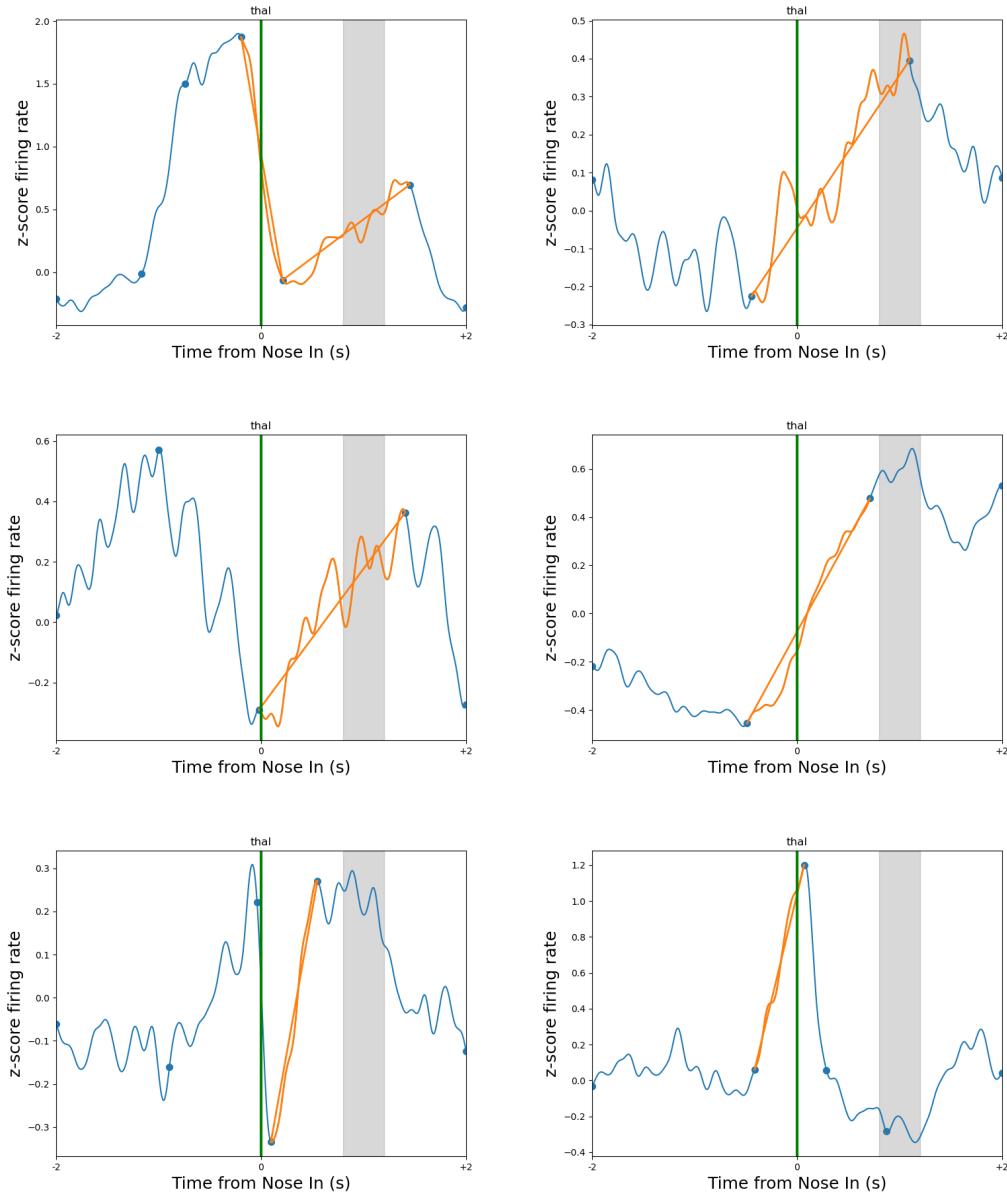


Figure B.5: Examples of ramp detection using the PELT algorithm for the thal. Orange section indicates detected ramp. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

B.2 Firing rates with no detected ramps aligned to Nose in

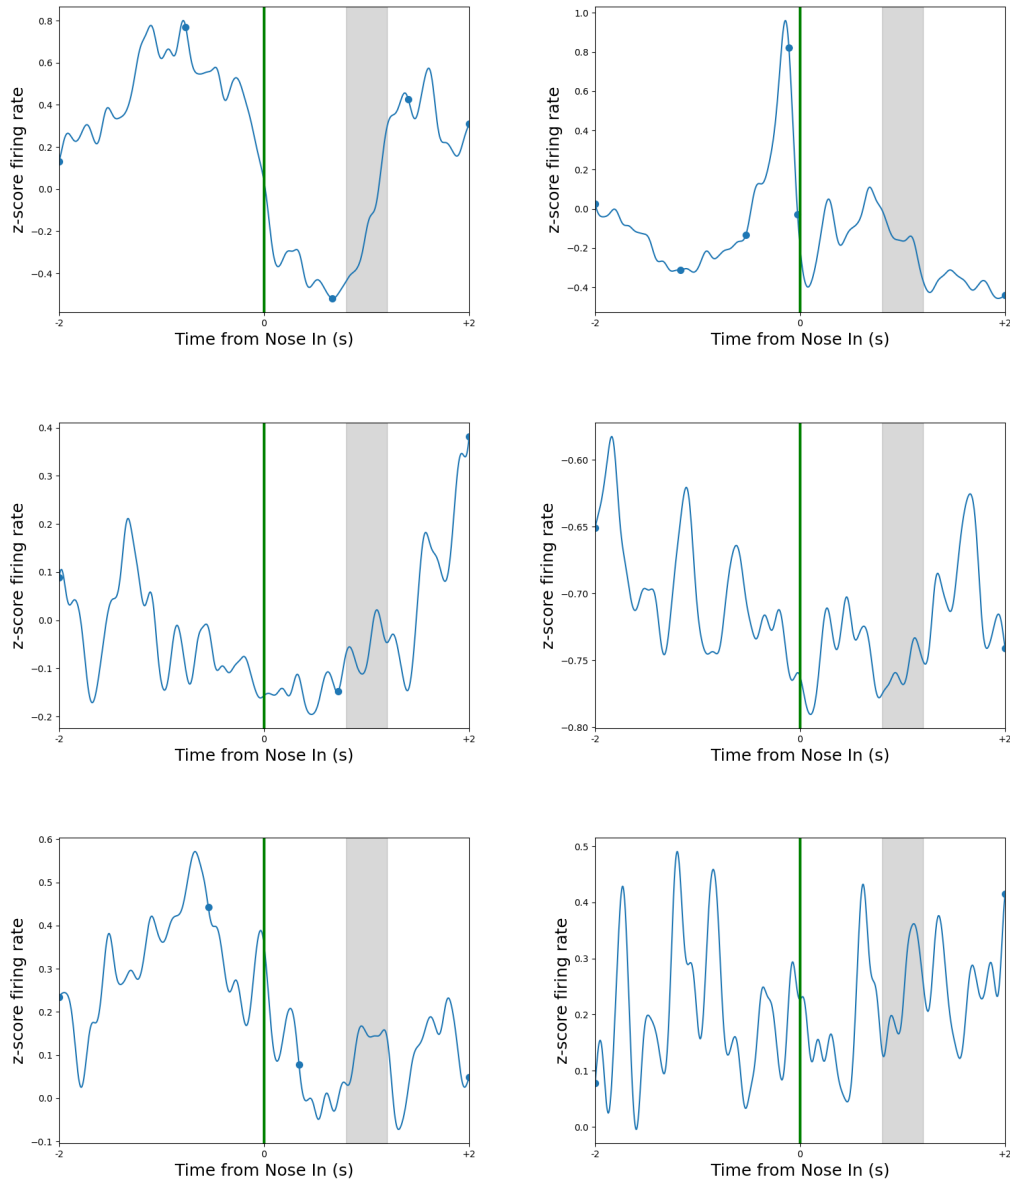


Figure B.6: Examples of firing rates with no ramps detected using the PELT algorithm for the Str. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

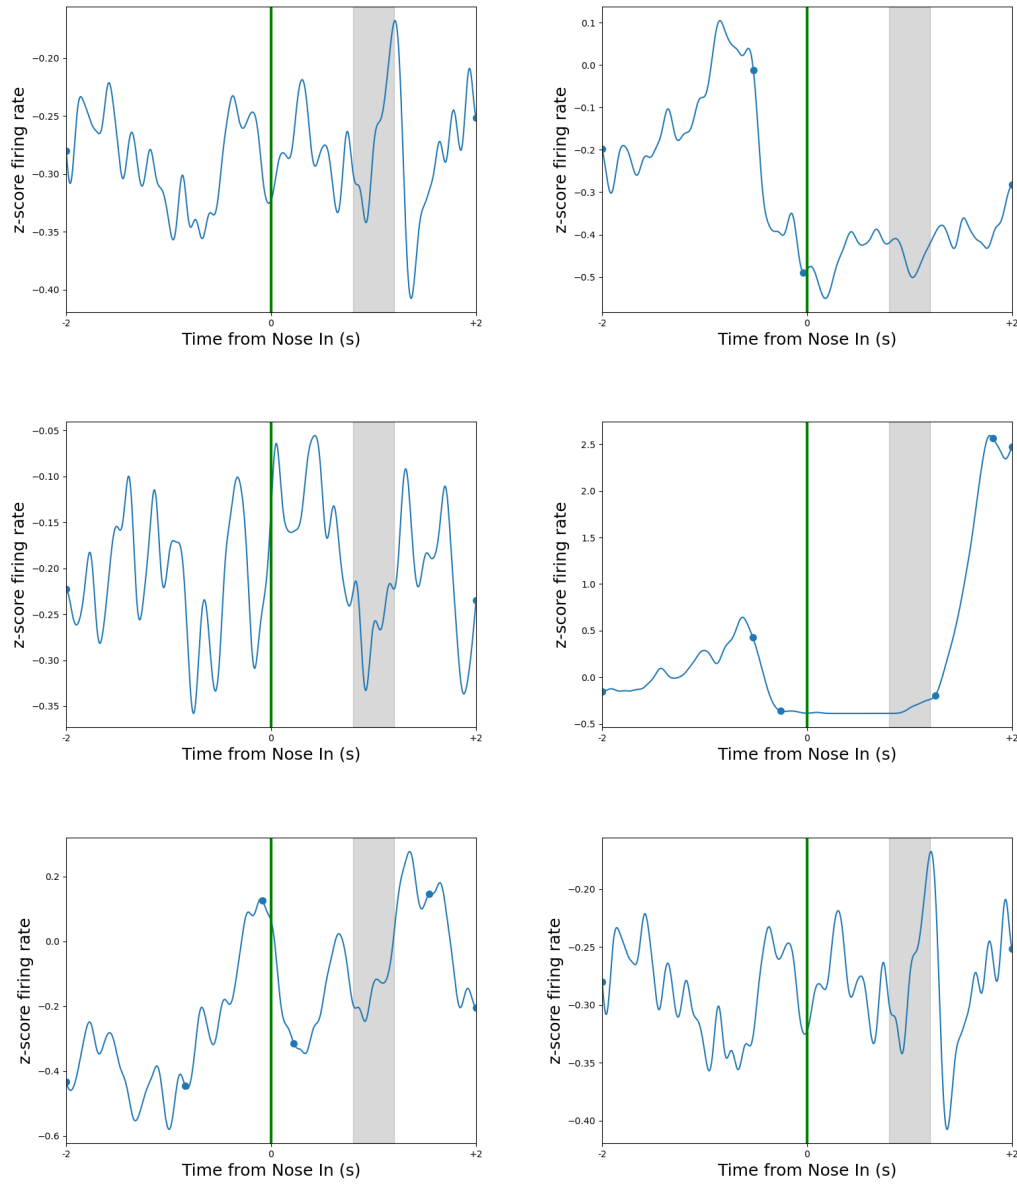


Figure B.7: Examples of firing rates with no ramps detected using the PELT algorithm for the GP. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

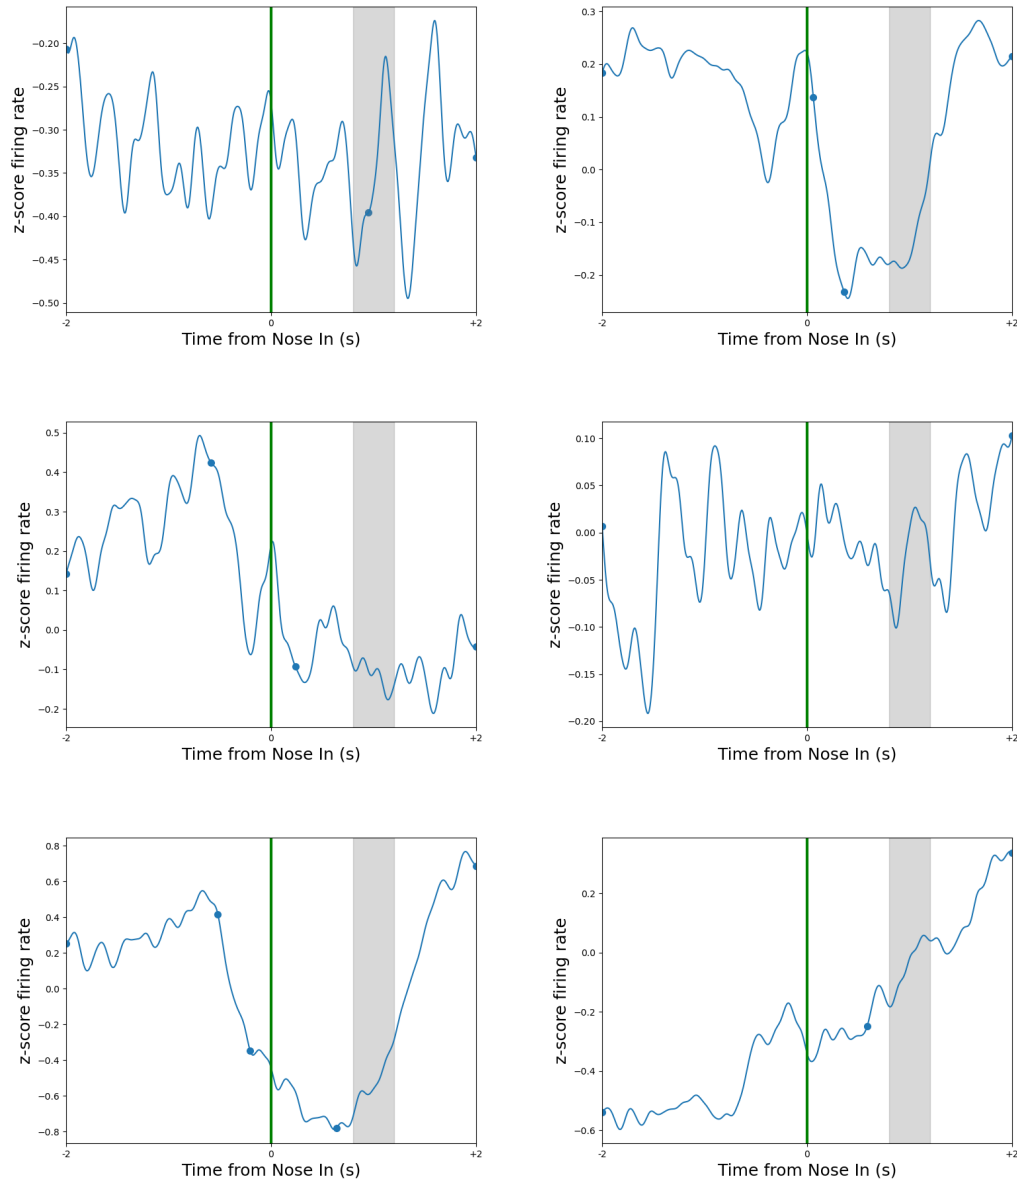


Figure B.8: Examples of firing rates with no ramps detected using the PELT algorithm for the SNr. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

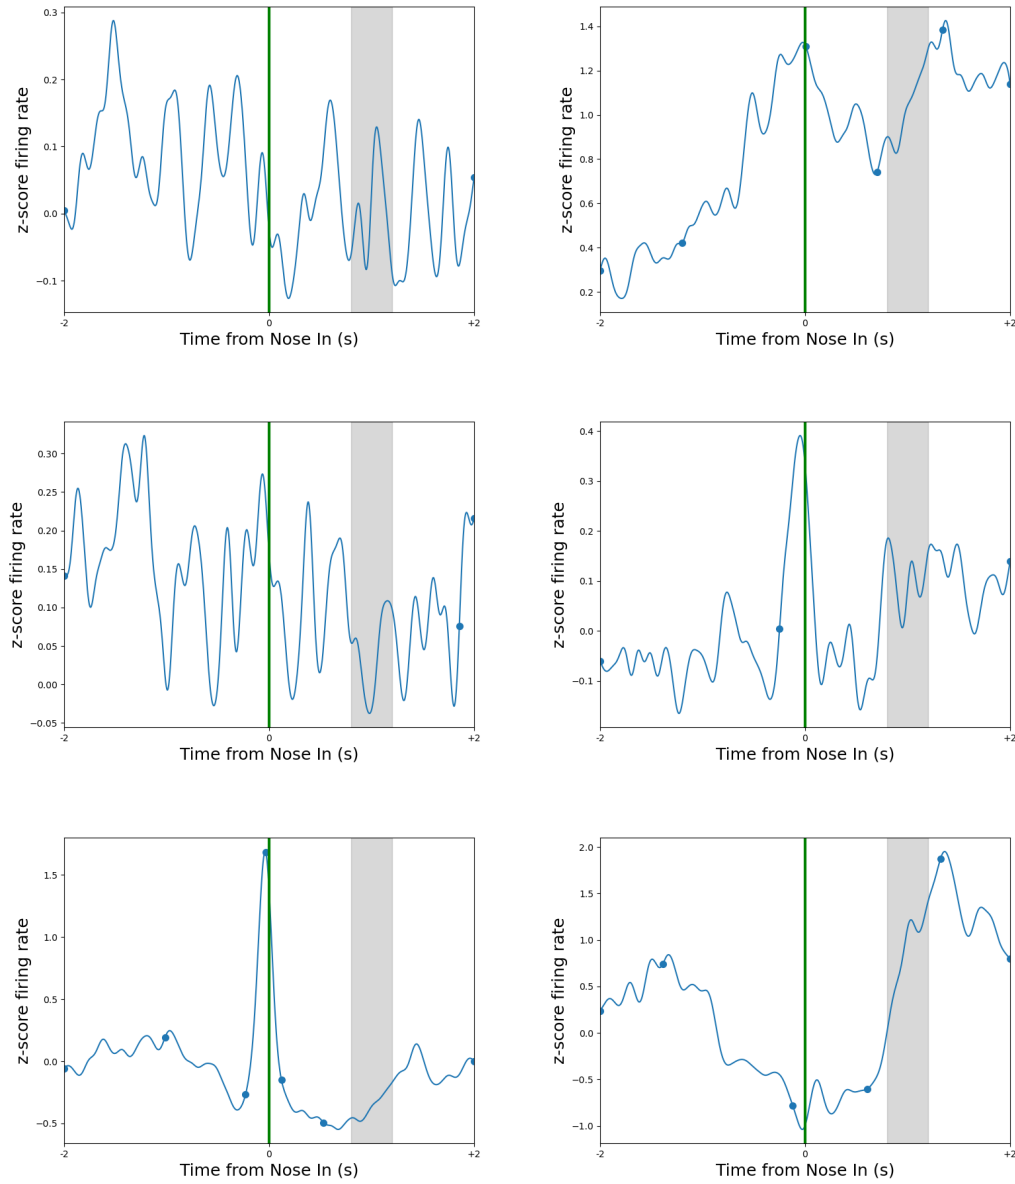


Figure B.9: Examples of firing rates with no ramps detected using the PELT algorithm for the STN. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

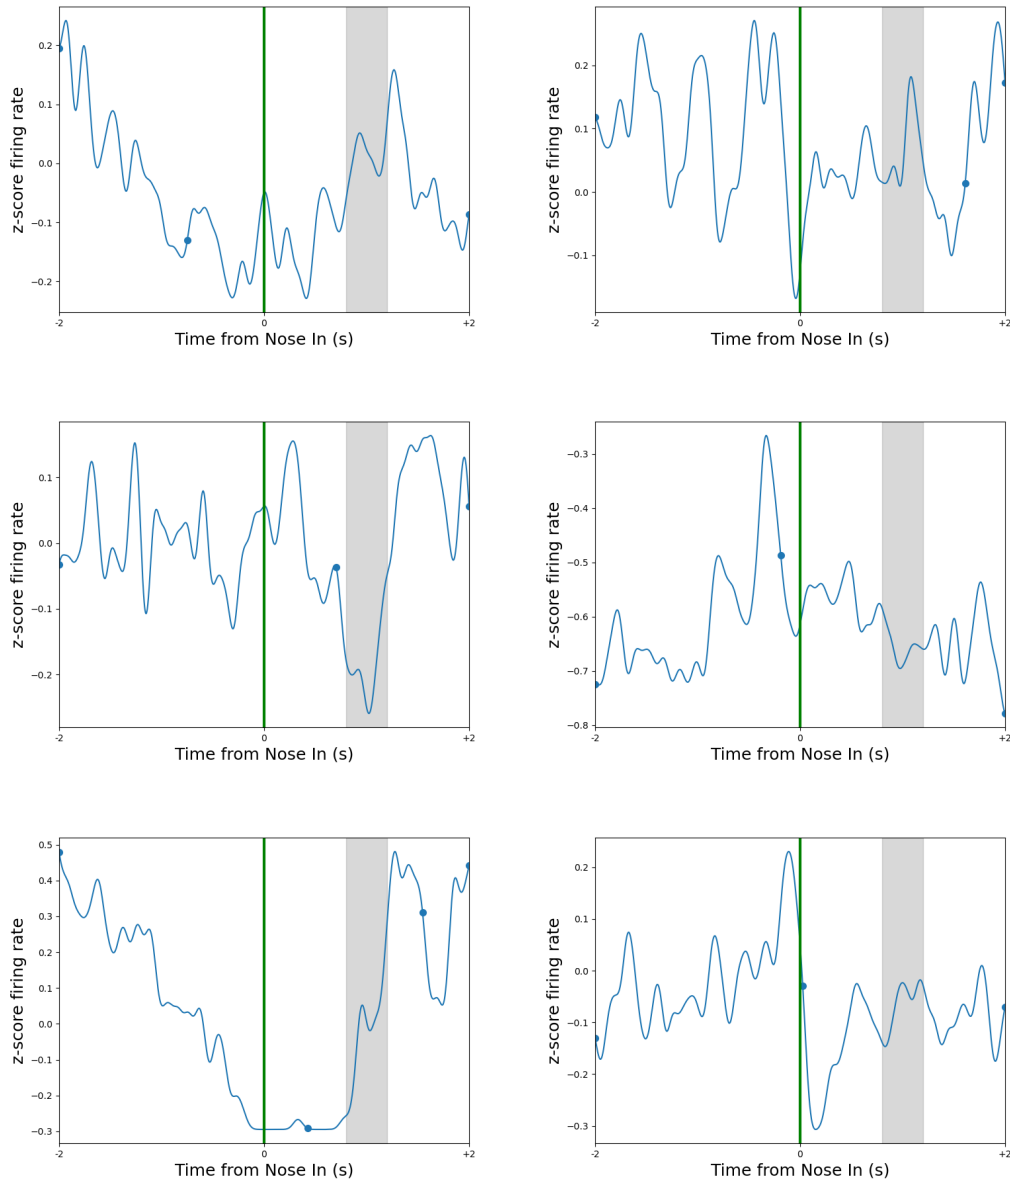


Figure B.10: Examples of firing rates with no ramps detected using the PELT algorithm for the thal. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

B.3 Detected ramps aligned to Go cue

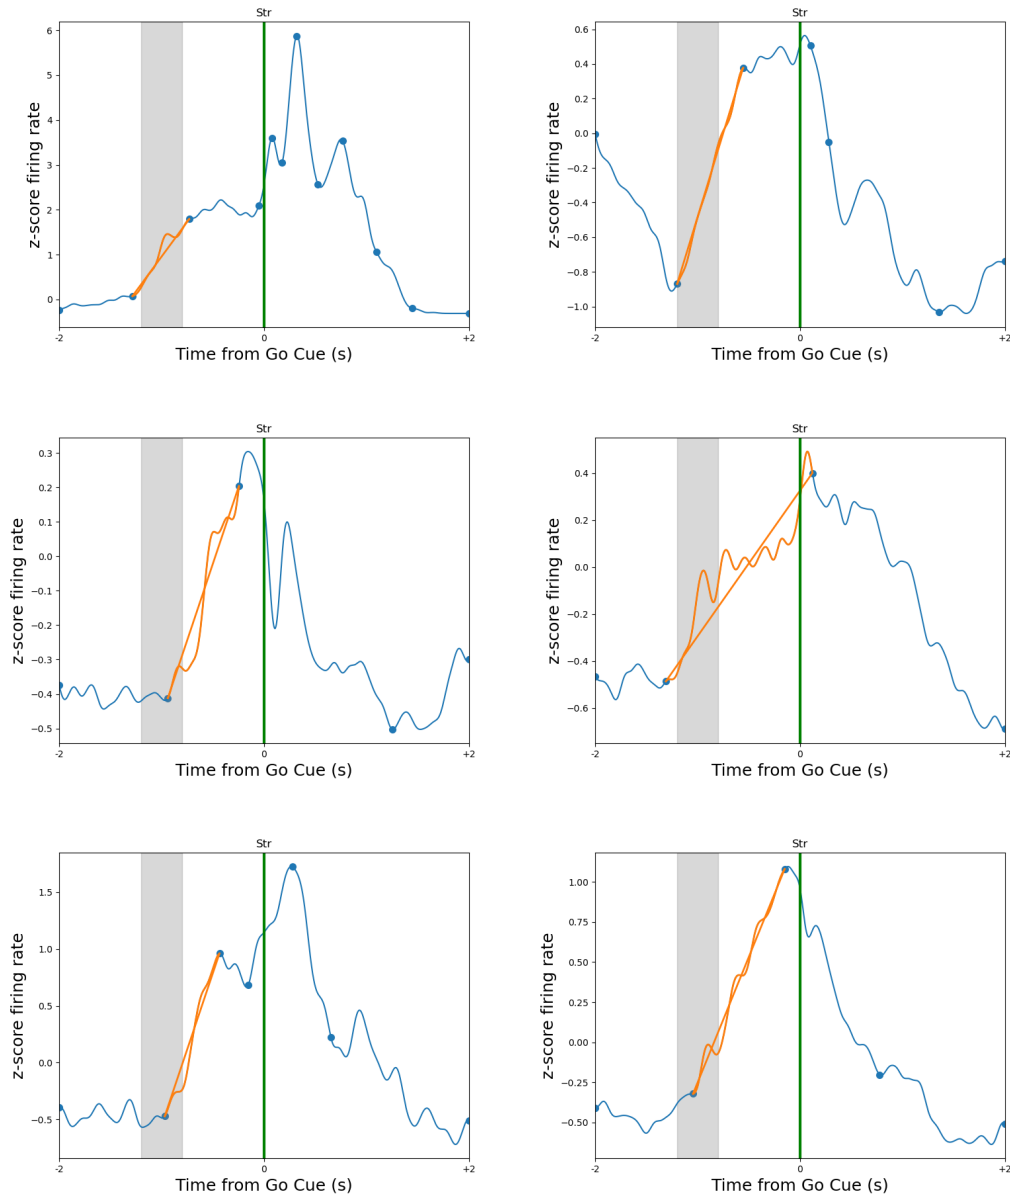


Figure B.11: Examples of ramp detection using the PELT algorithm for the Str. Orange section indicates detected ramp. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

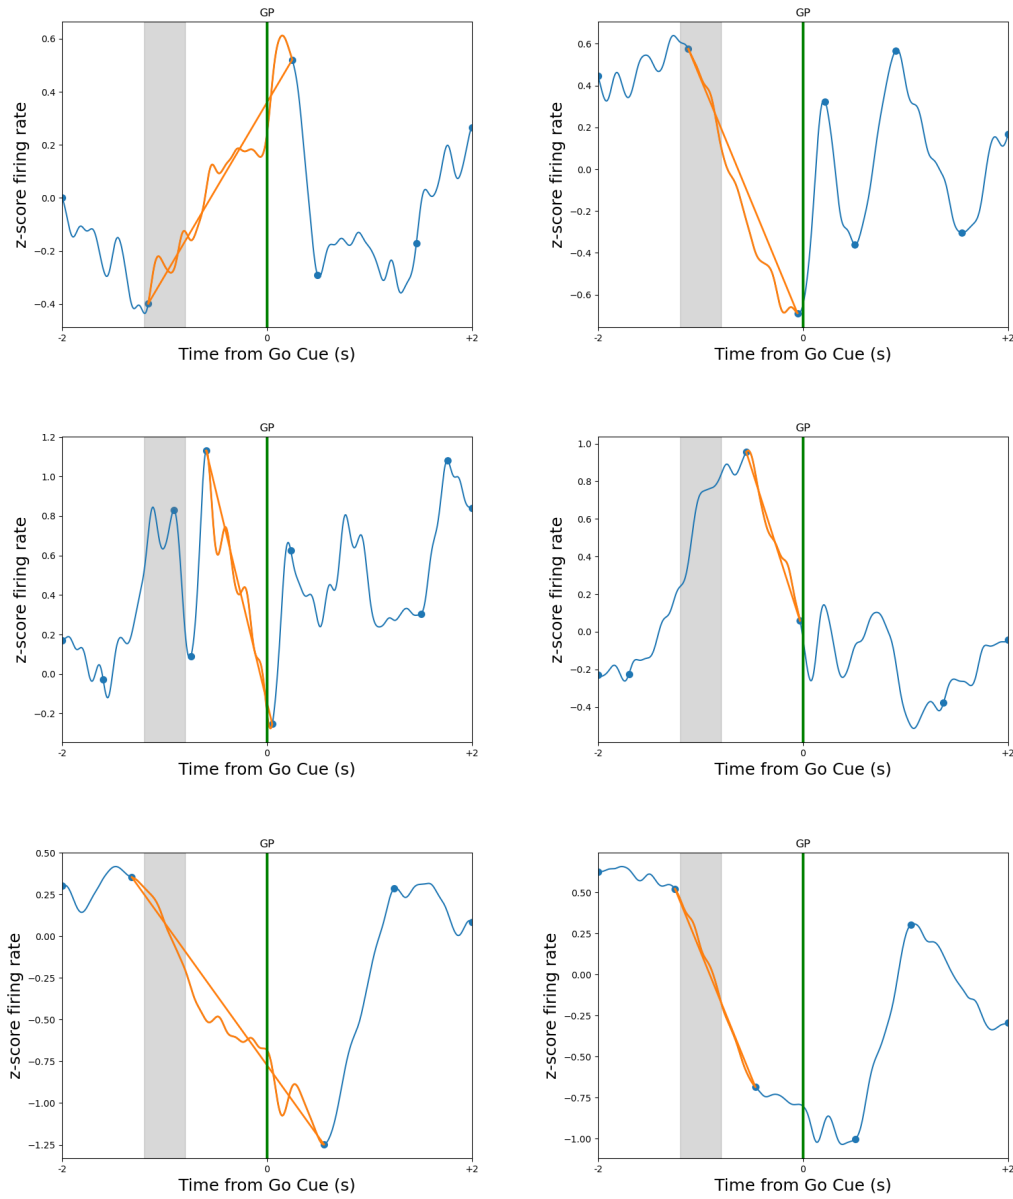


Figure B.12: Examples of ramp detection using the PELT algorithm for the GP. Orange section indicates detected ramp. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

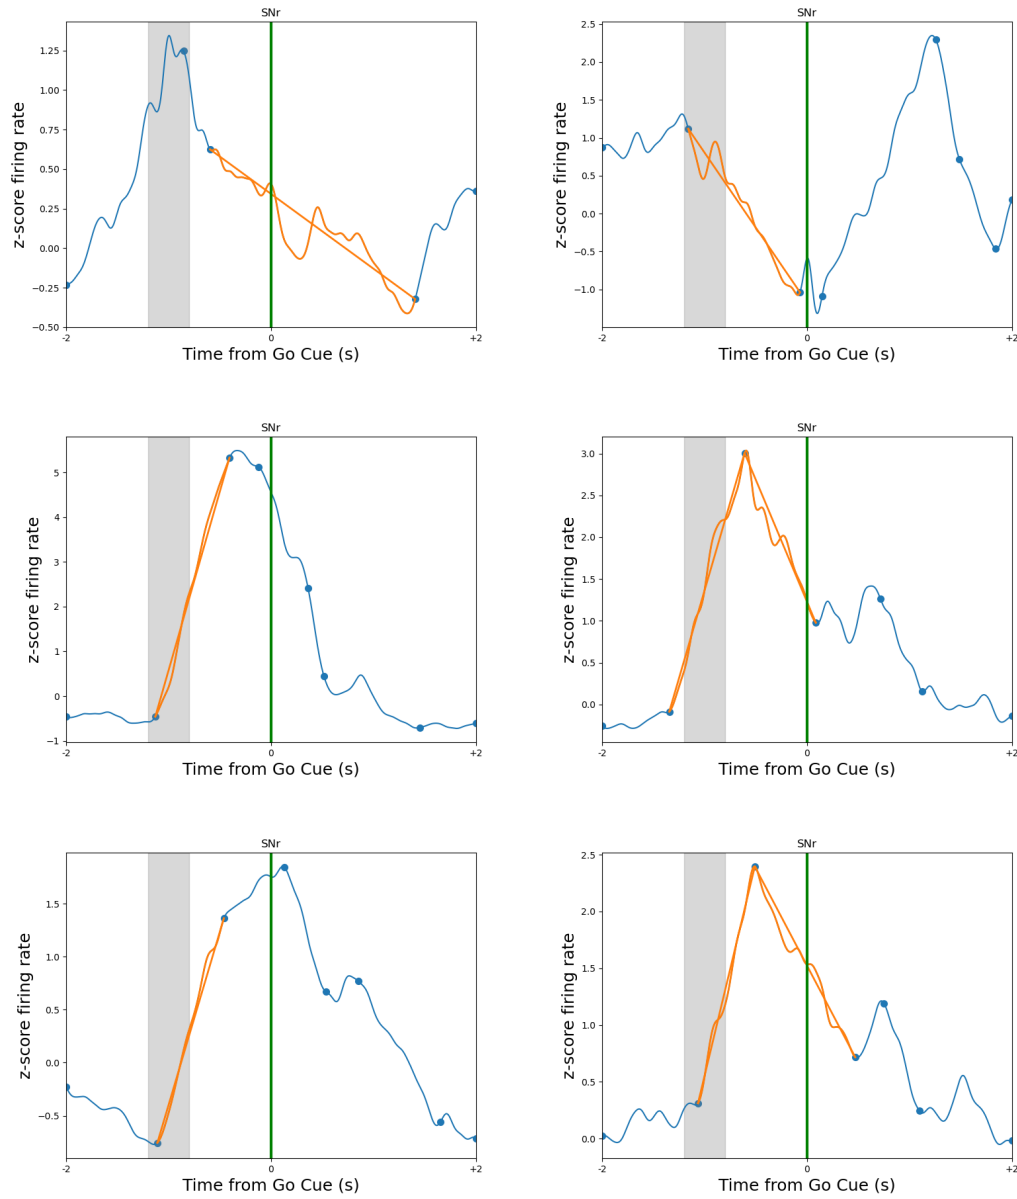


Figure B.13: Examples of ramp detection using the PELT algorithm for the SNr. Orange section indicates detected ramp. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

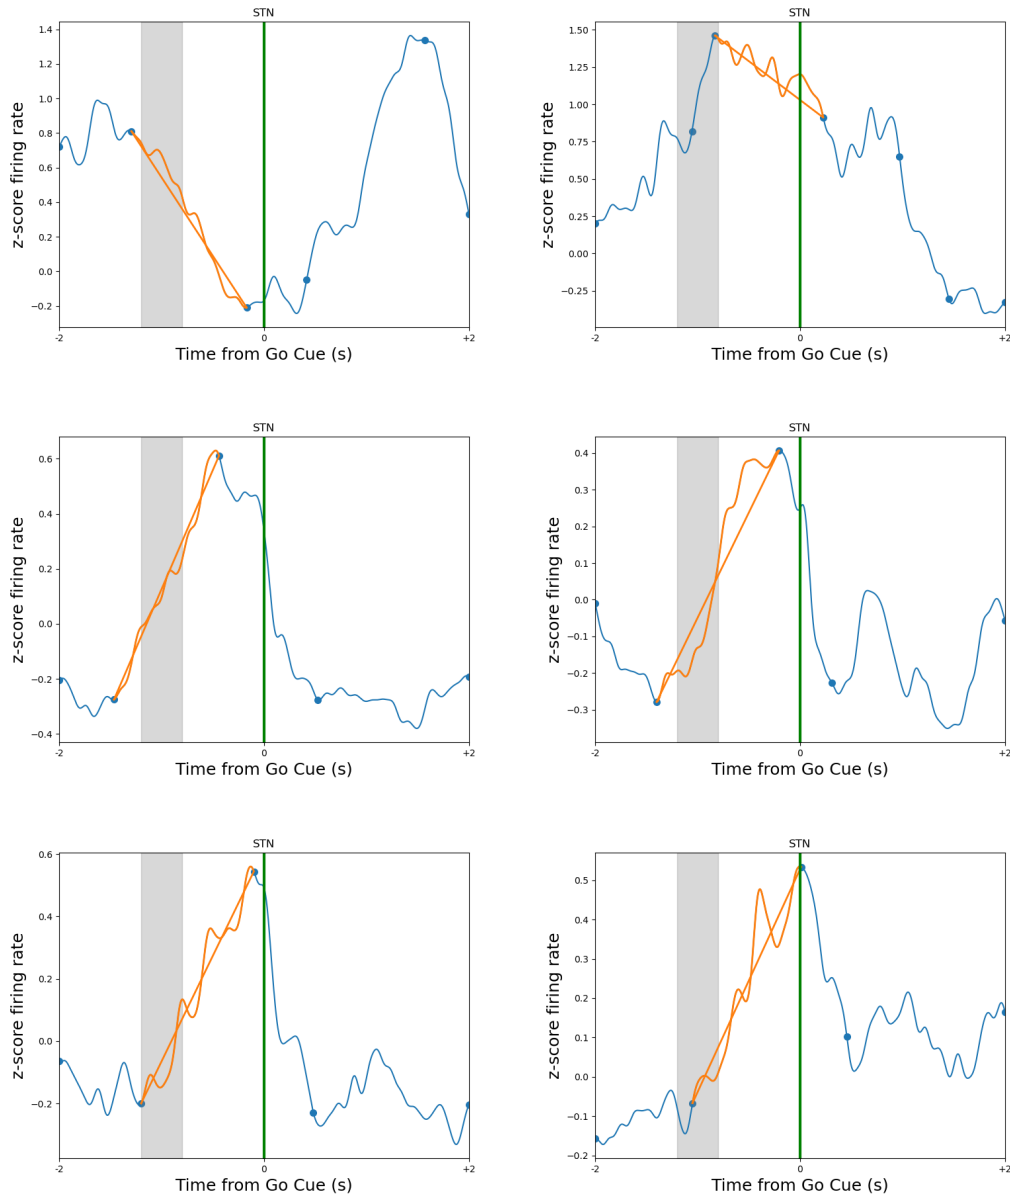


Figure B.14: Examples of ramp detection using the PELT algorithm for the STN. Orange section indicates detected ramp. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

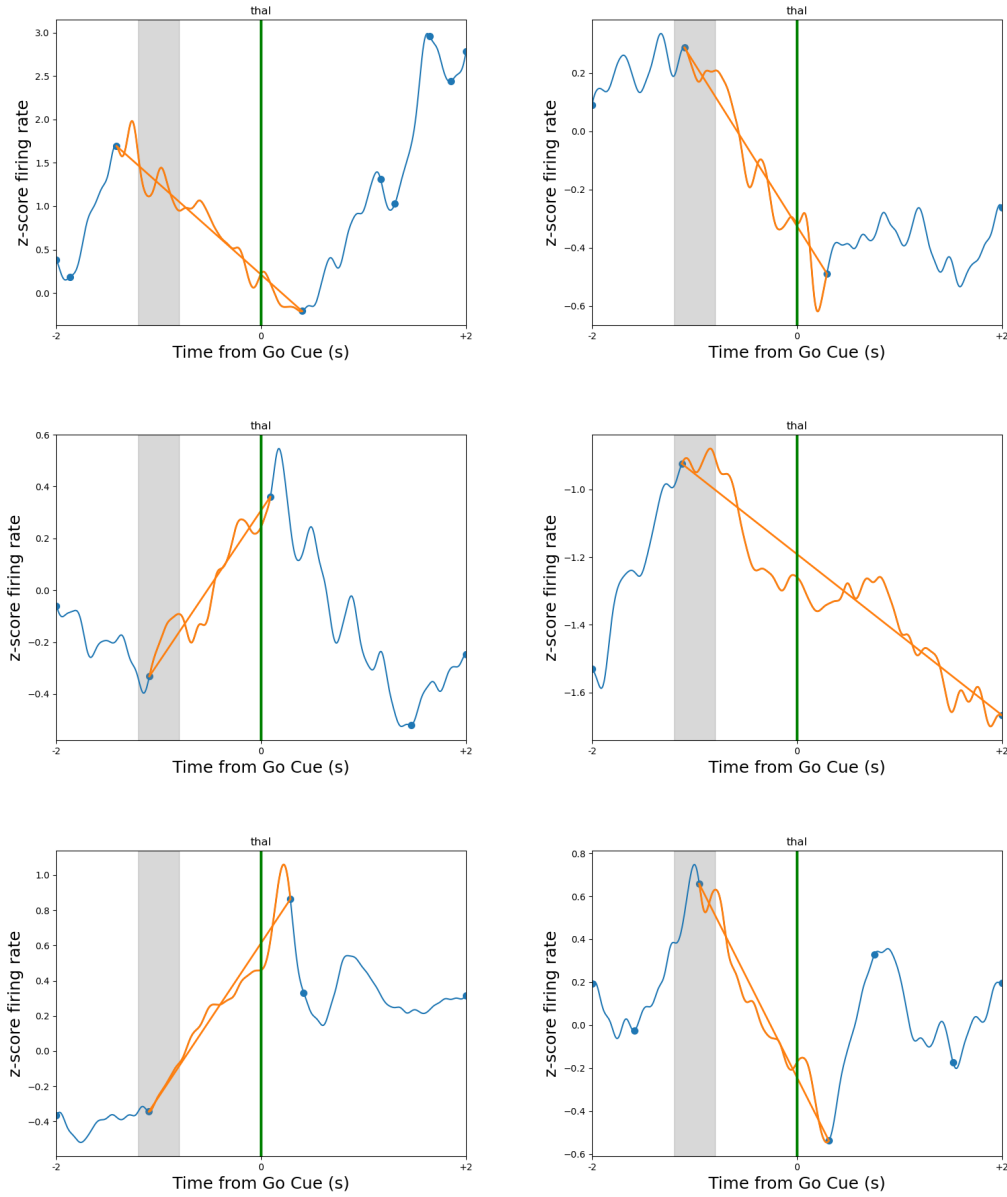


Figure B.15: Examples of ramp detection using the PELT algorithm for the thal. Orange section indicates detected ramp. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

B.4 Firing rates with no detected ramps aligned to Go cue

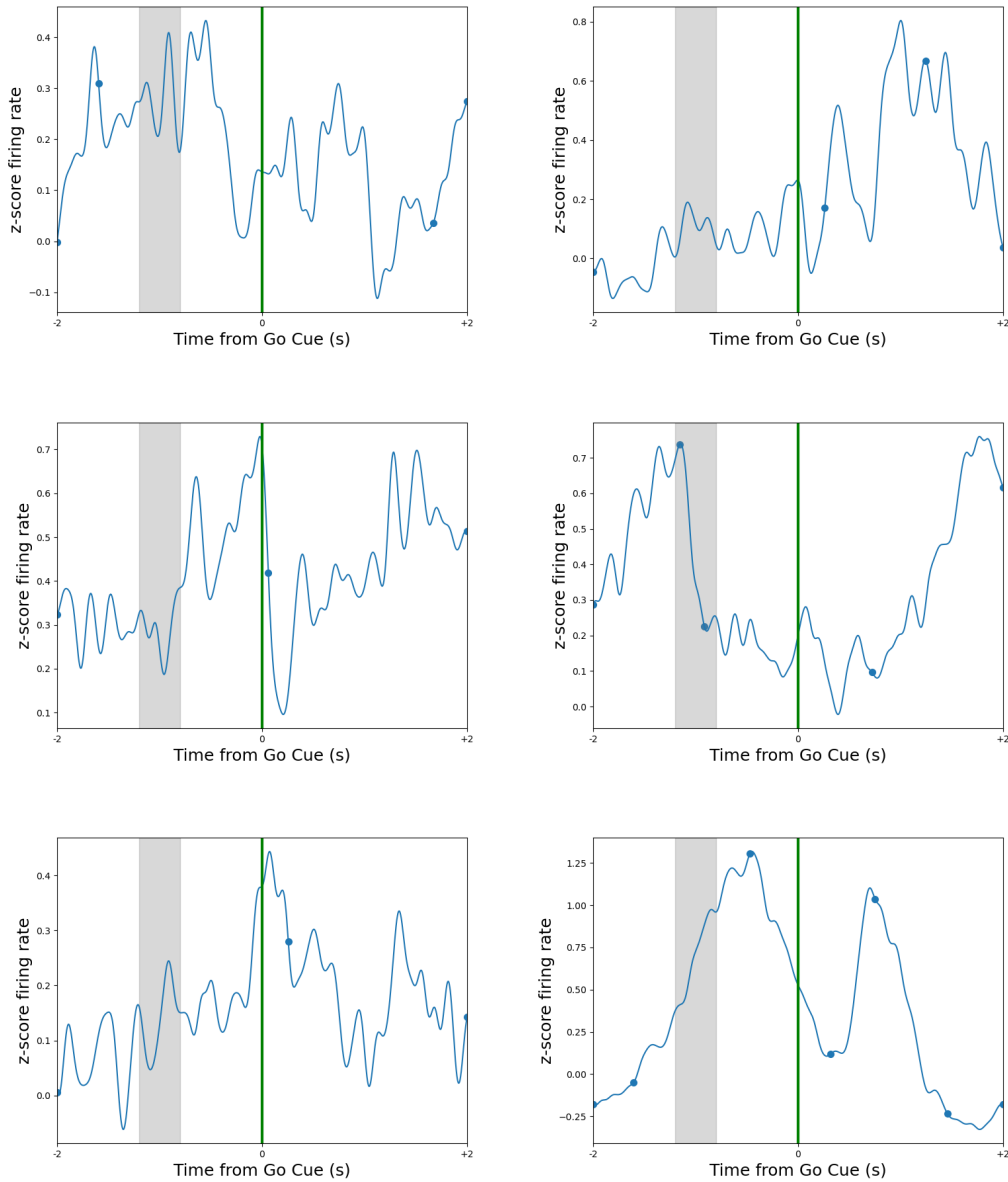


Figure B.16: Examples of firing rates with no ramps detected using the PELT algorithm for the Str. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

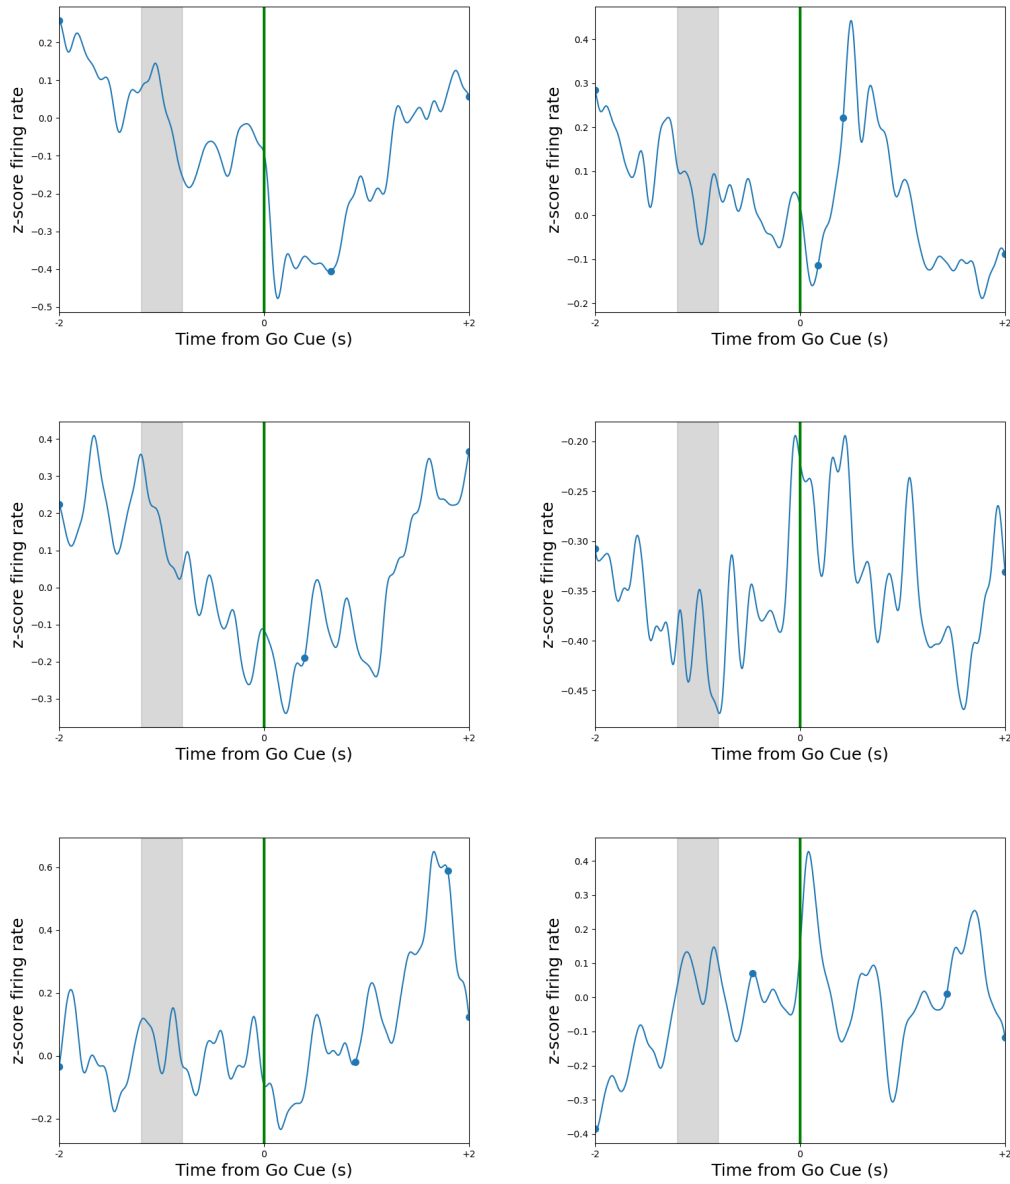


Figure B.17: Examples of firing rates with no ramps detected using the PELT algorithm for the GP. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

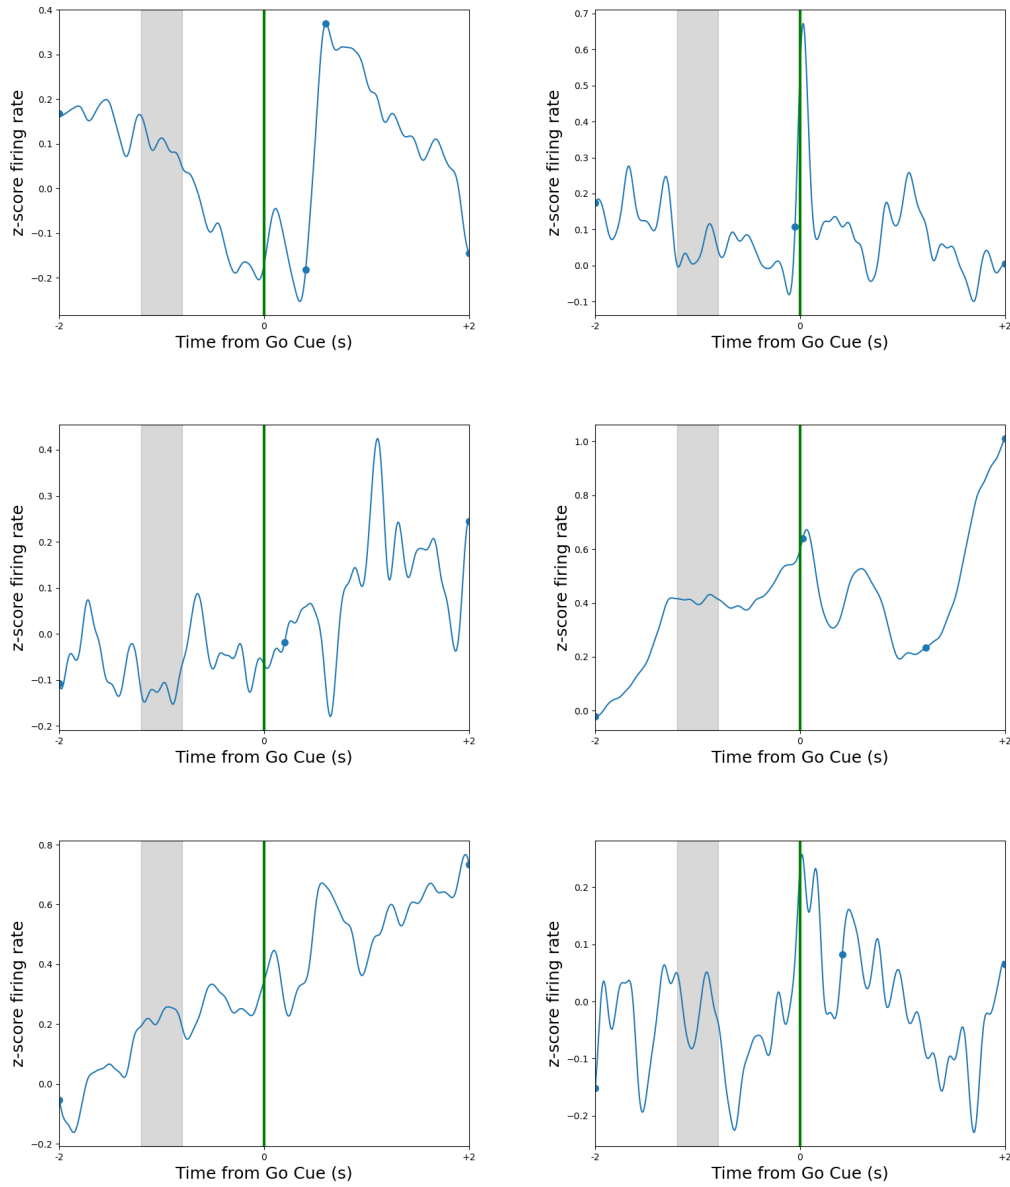


Figure B.18: Examples of firing rates with no ramps detected using the PELT algorithm for the SNr. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

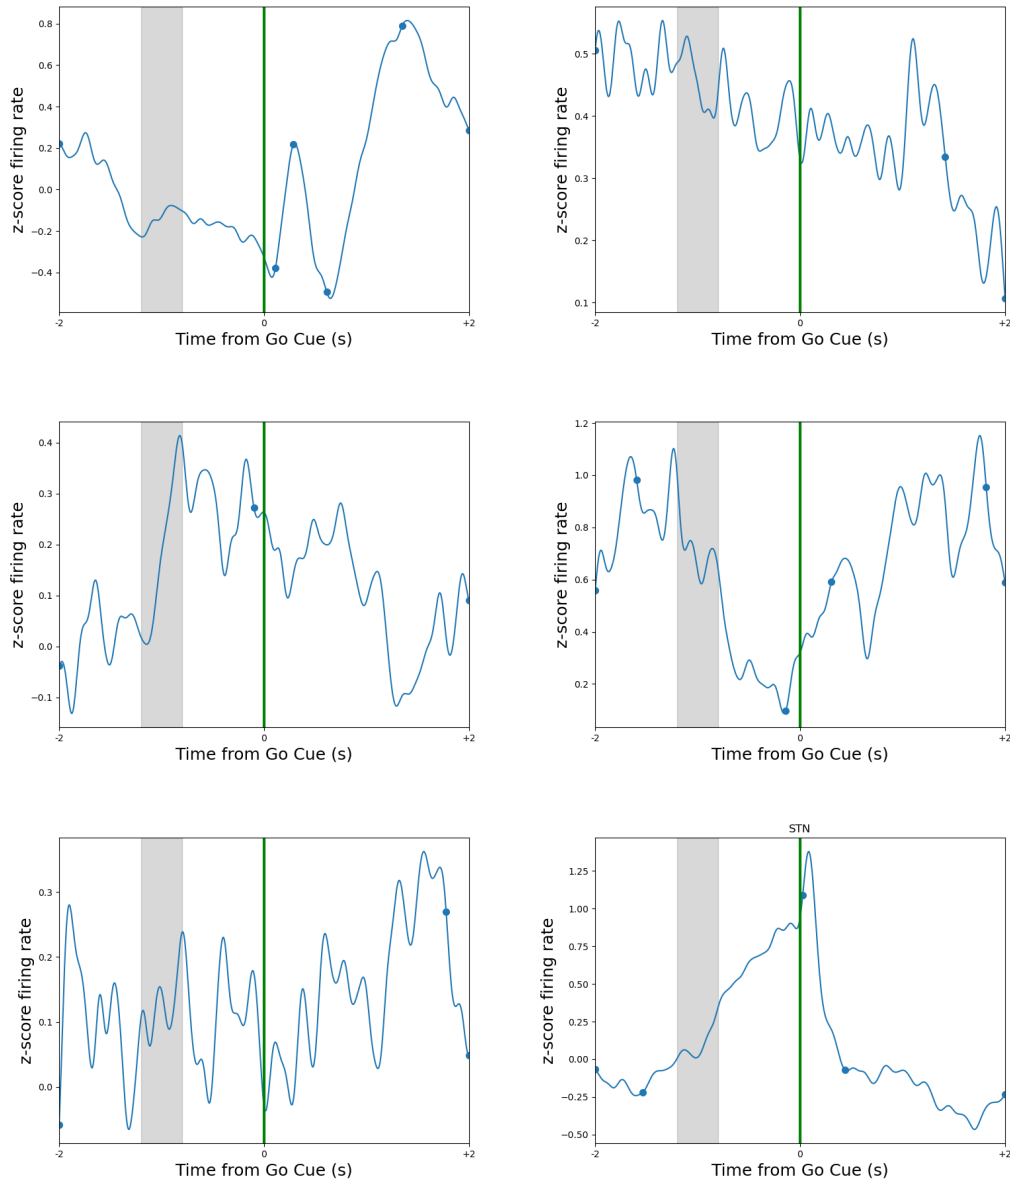


Figure B.19: Examples of firing rates with no ramps detected using the PELT algorithm for the STN. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

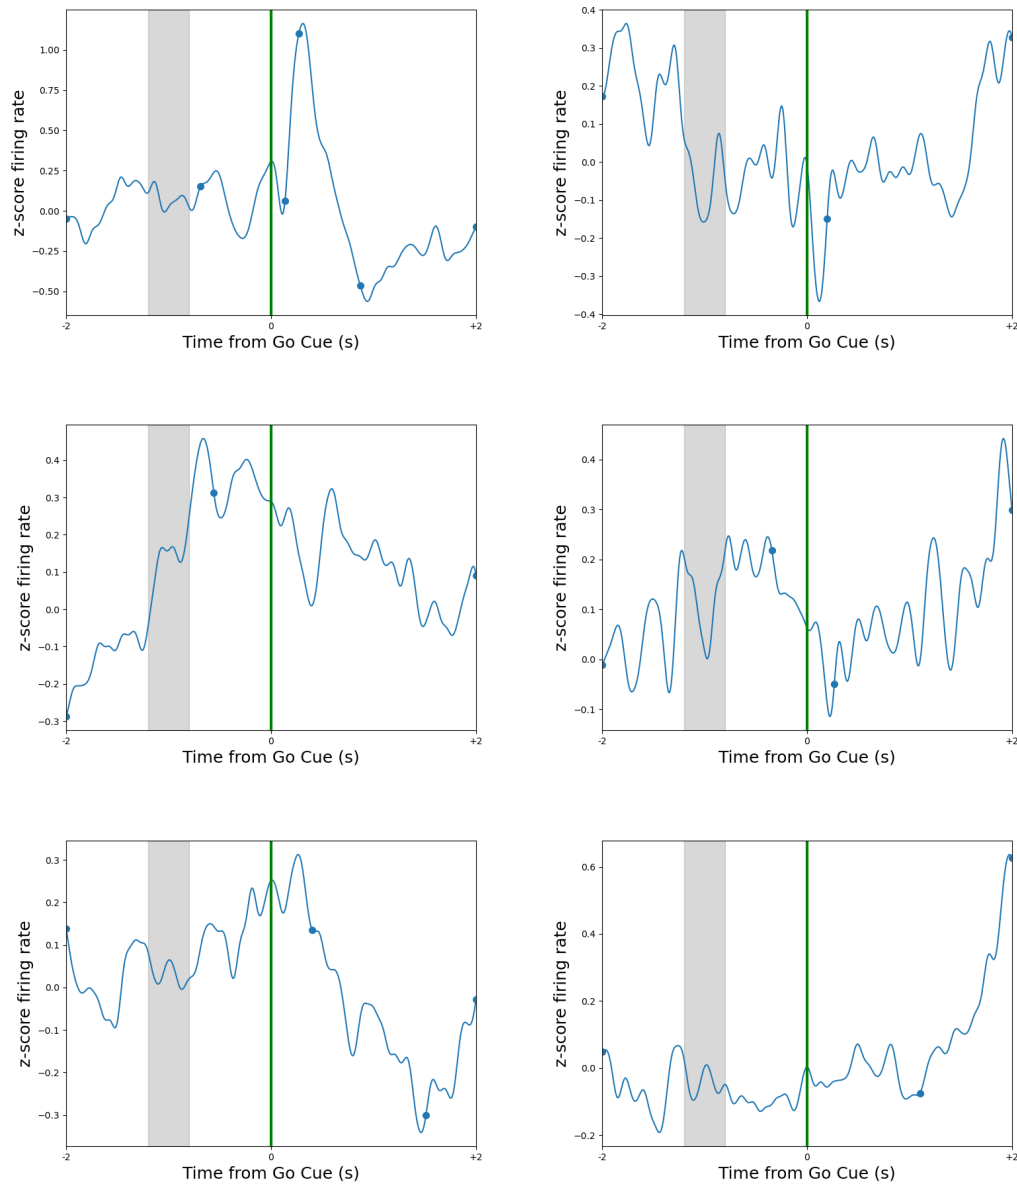


Figure B.20: Examples of firing rates with no ramps detected using the PELT algorithm for the thal. Only ramps that start within 0.5 seconds of Nose in and with a quality value of more than 50 are included.

B.5 Single trial analysis of positive detected ramps.

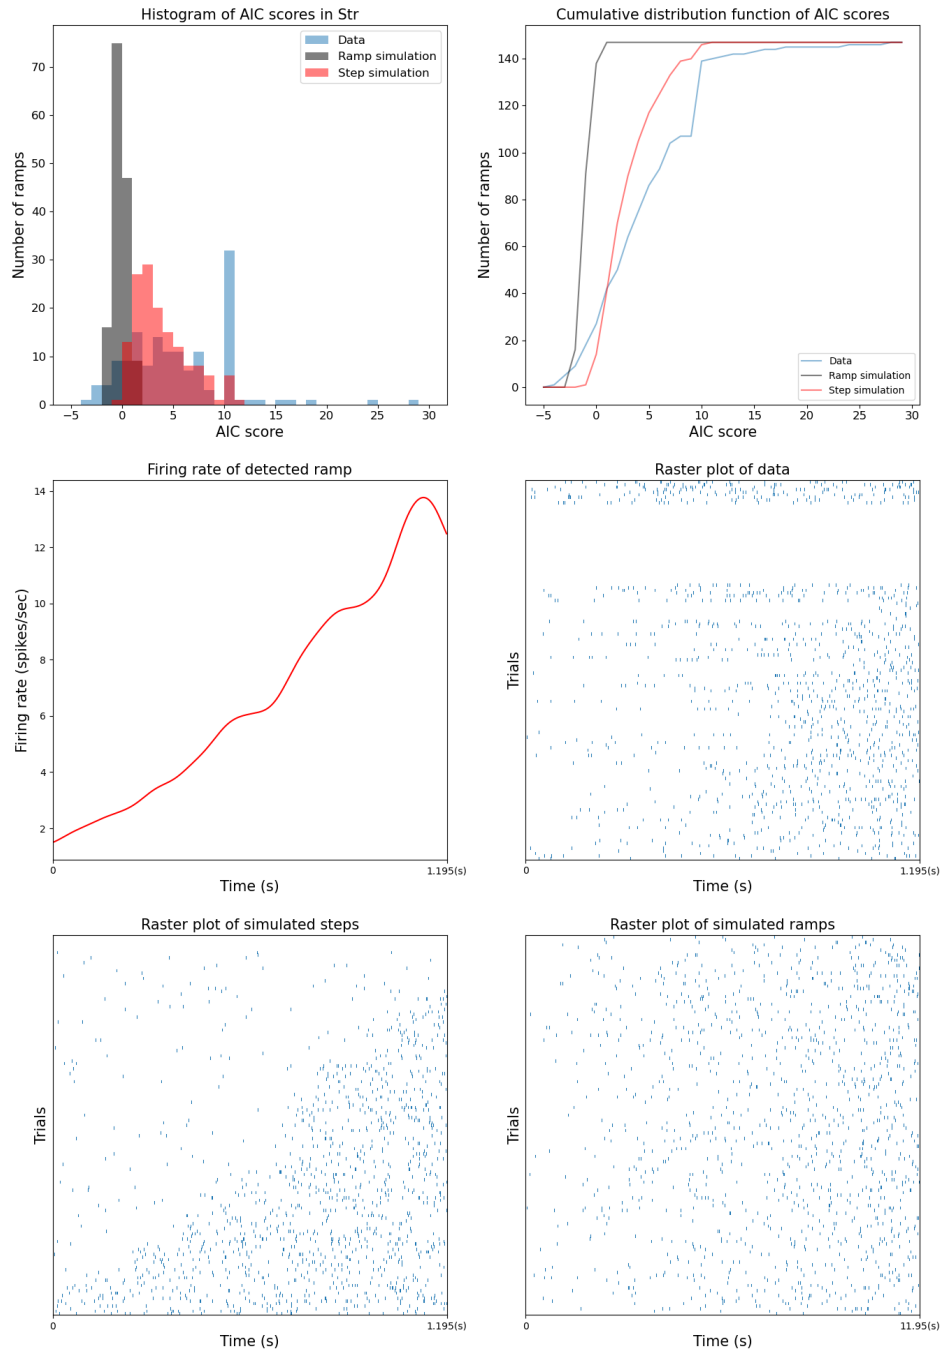


Figure B.21: Detected ramp in Str has ramping single-trial firing rates. Top row shows AIC score distribution on the left and the cumulative distribution of AIC scores on the right. Negative scores suggest the trial is more likely to be ramping, while positive scores suggest the trial is more likely to be stepping. Middle row shows the firing rate of the detected ramp on the left with the raster plot of the trials on the right. Bottom row shows raster plots of simulated steps and ramps.

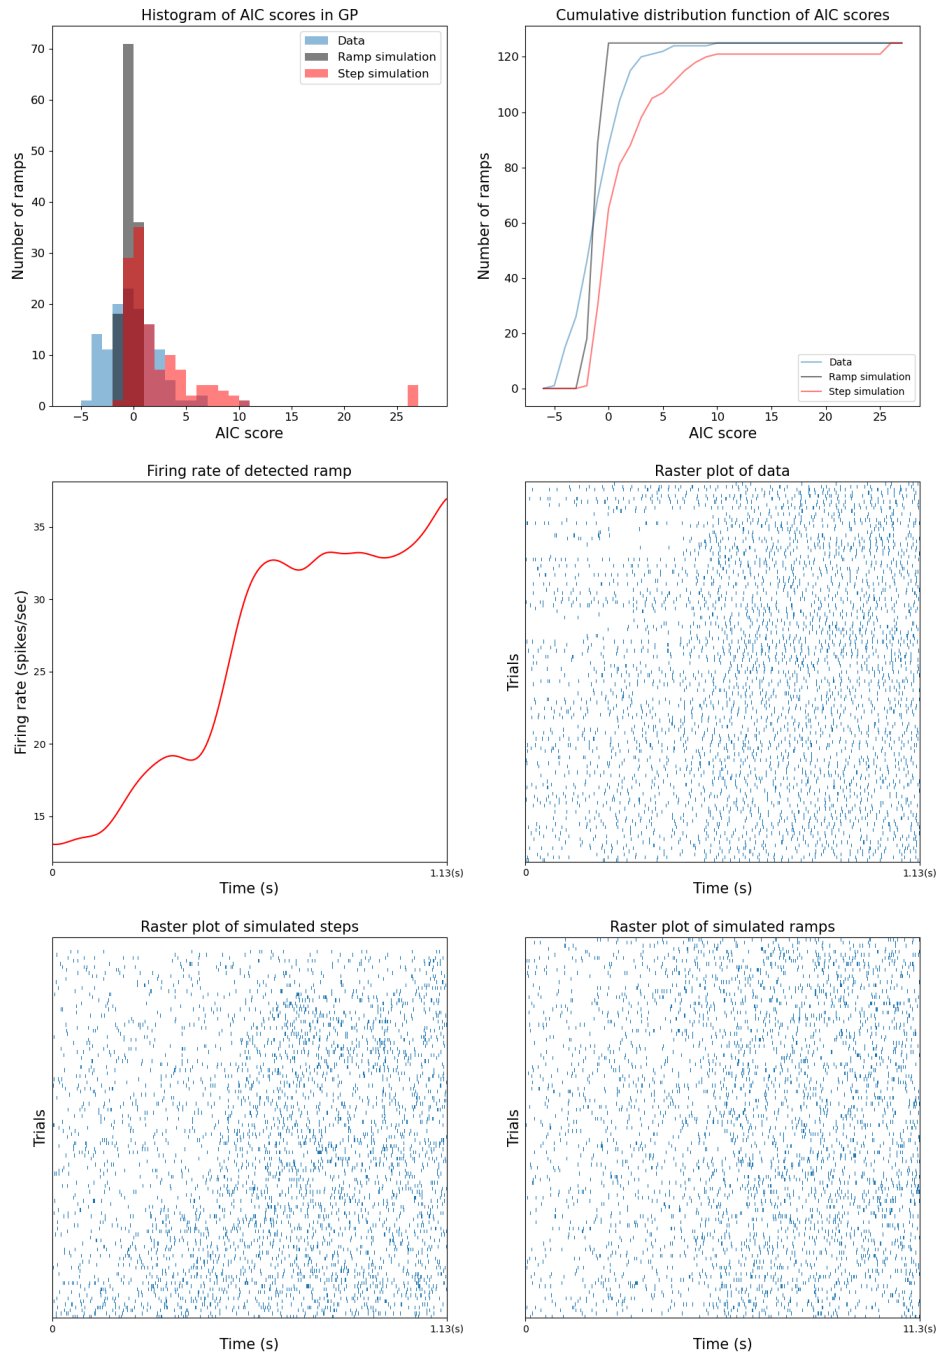


Figure B.22: Detected ramp in GP has ramping single-trial firing rates. Top row shows AIC score distribution on the left and the cumulative distribution of AIC scores on the right. Negative scores suggest the trial is more likely to be ramping while positive scores suggest the trial is more likely to be stepping. Middle row shows the firing rate of the detected ramp on the left with the raster plot of the trials on the right. Bottom row shows raster plots of simulated steps and ramps.

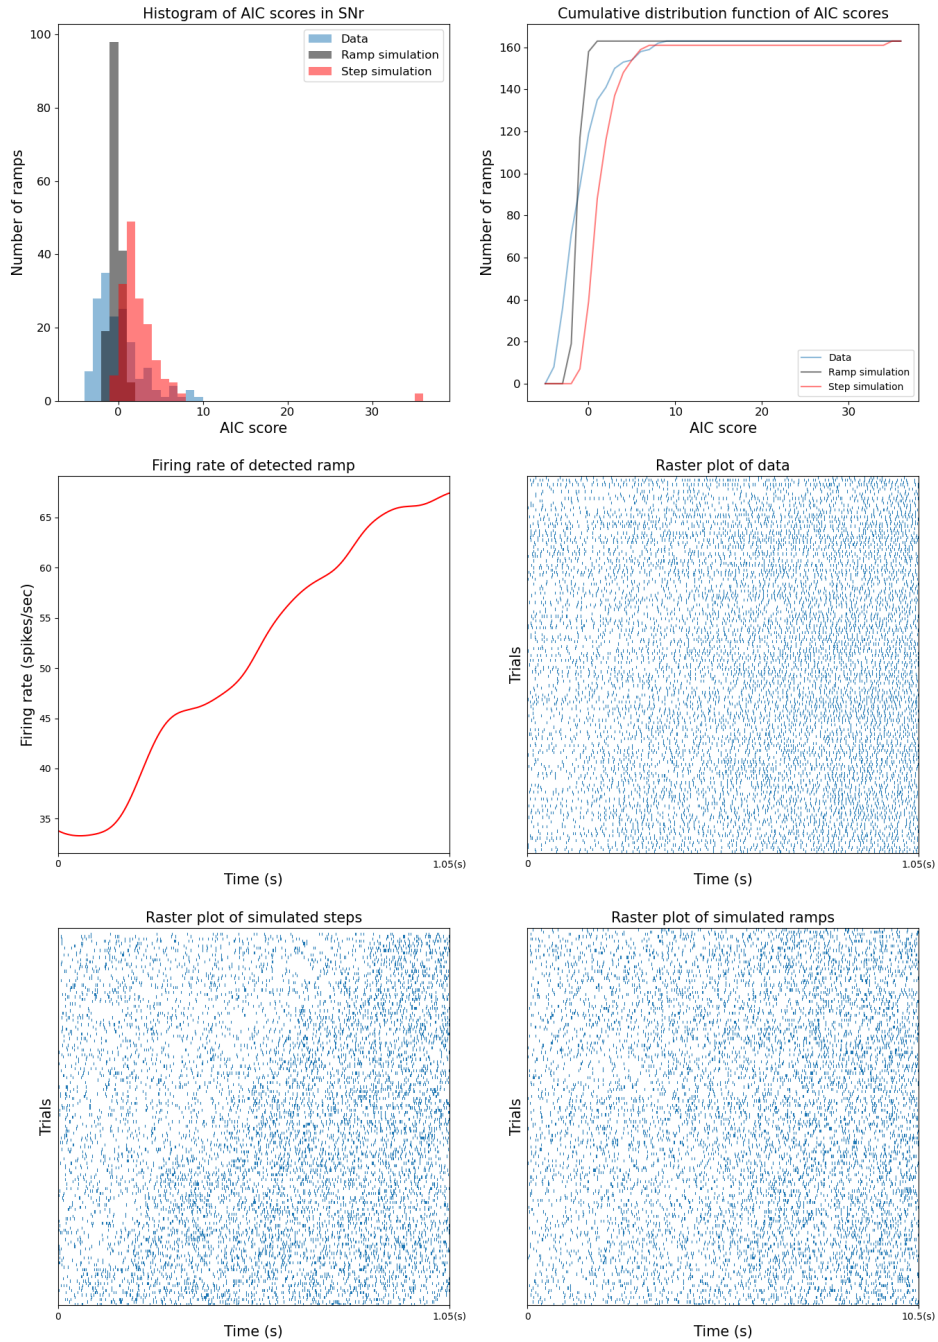


Figure B.23: Detected ramp in SNr has ramping single-trial firing rates. Top row shows AIC score distribution on the left and the cumulative distribution of AIC scores on the right. Negative scores suggest the trial is more likely to be ramping while positive scores suggest the trial is more likely to be stepping. Middle row shows the firing rate of the detected ramp on the left with the raster plot of the trials on the right. Bottom row shows raster plots of simulated steps and ramps.

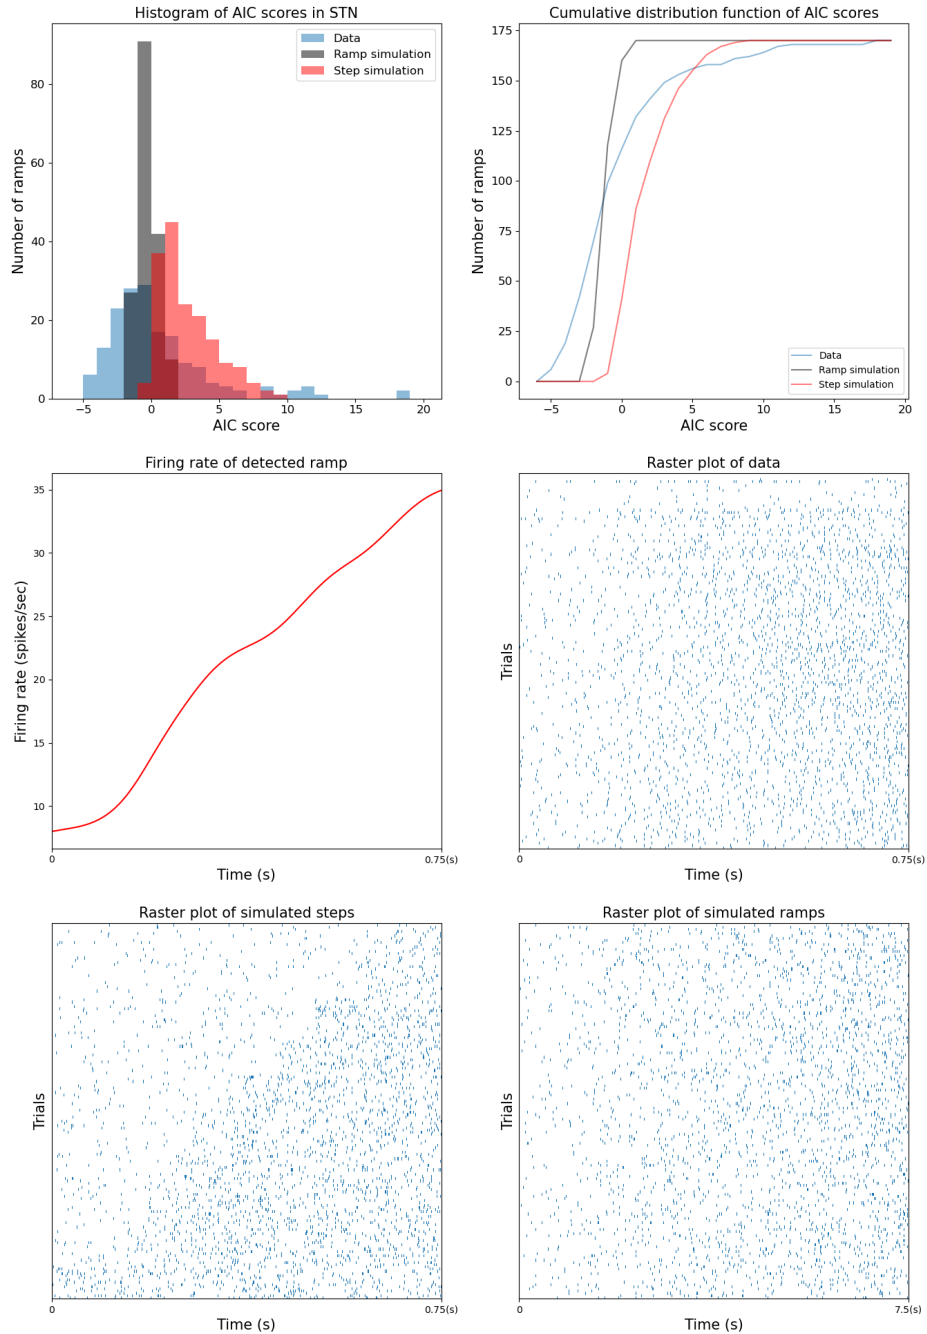


Figure B.24: Detected ramp in STN has ramping single-trial firing rates. Top row shows AIC score distribution on the left and the cumulative distribution of AIC scores on the right. Negative scores suggest the trial is more likely to be ramping while positive scores suggest the trial is more likely to be stepping. Middle row shows the firing rate of the detected ramp on the left with the raster plot of the trials on the right. Bottom row shows raster plots of simulated steps and ramps.

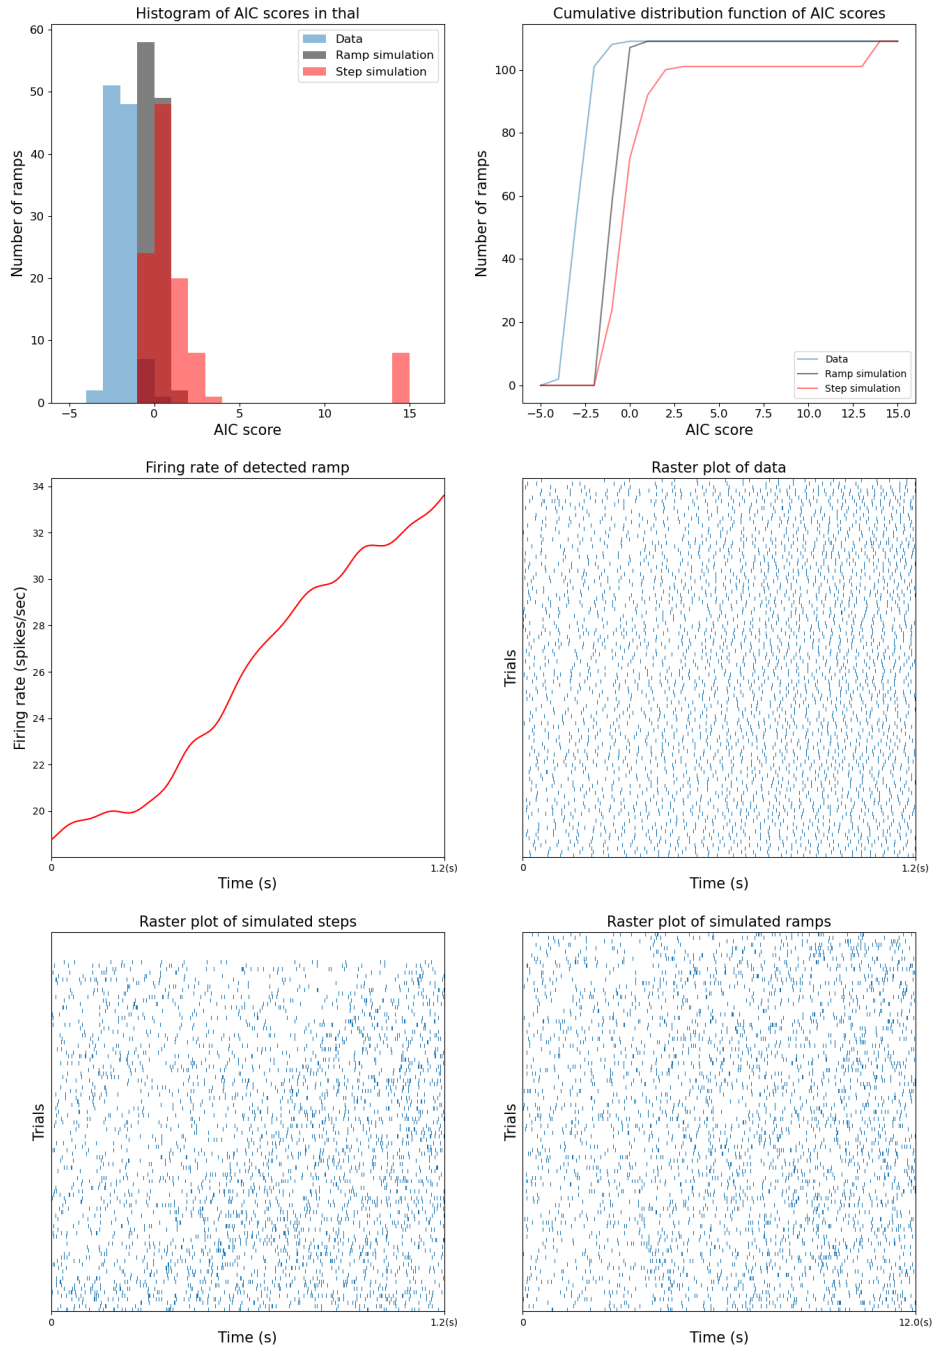


Figure B.25: Detected ramp in thal has ramping single-trial firing rates. Top row shows AIC score distribution on the left and the cumulative distribution of AIC scores on the right. Negative scores suggest the trial is more likely to be ramping while positive scores suggest the trial is more likely to be stepping. Middle row shows the firing rate of the detected ramp on the left with the raster plot of the trials on the right. Bottom row shows raster plots of simulated steps and ramps.

B.6 Single trial analysis of negative detected ramps.

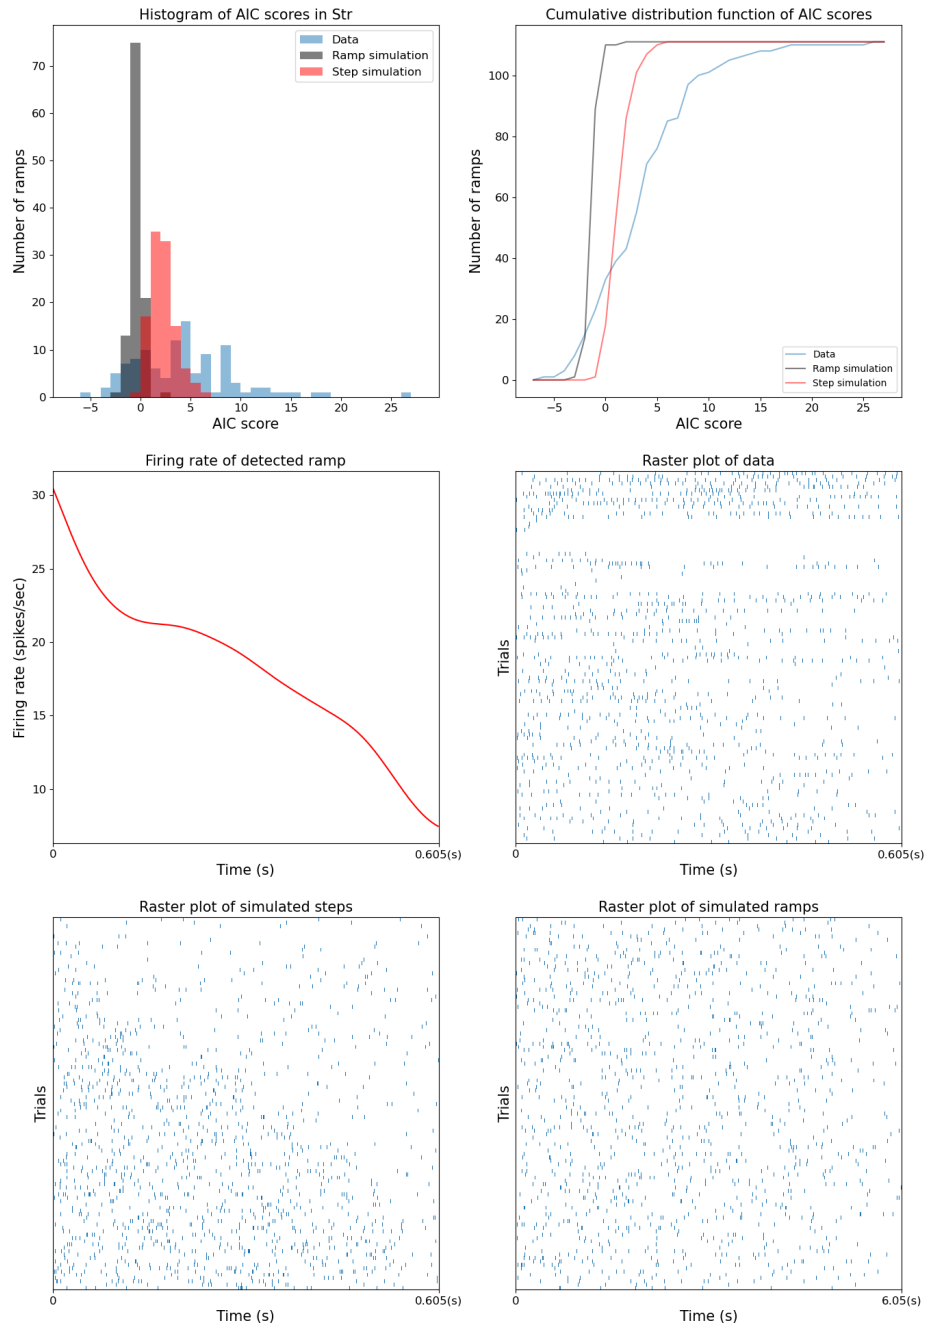


Figure B.26: Detected ramp in Str has ramping single-trial firing rates. Top row shows AIC score distribution on the left and the cumulative distribution of AIC scores on the right. Negative scores suggest the trial is more likely to be ramping while positive scores suggest the trial is more likely to be stepping. Middle row shows the firing rate of the detected ramp on the left with the raster plot of the trials on the right. Bottom row shows raster plots of simulated steps and ramps.

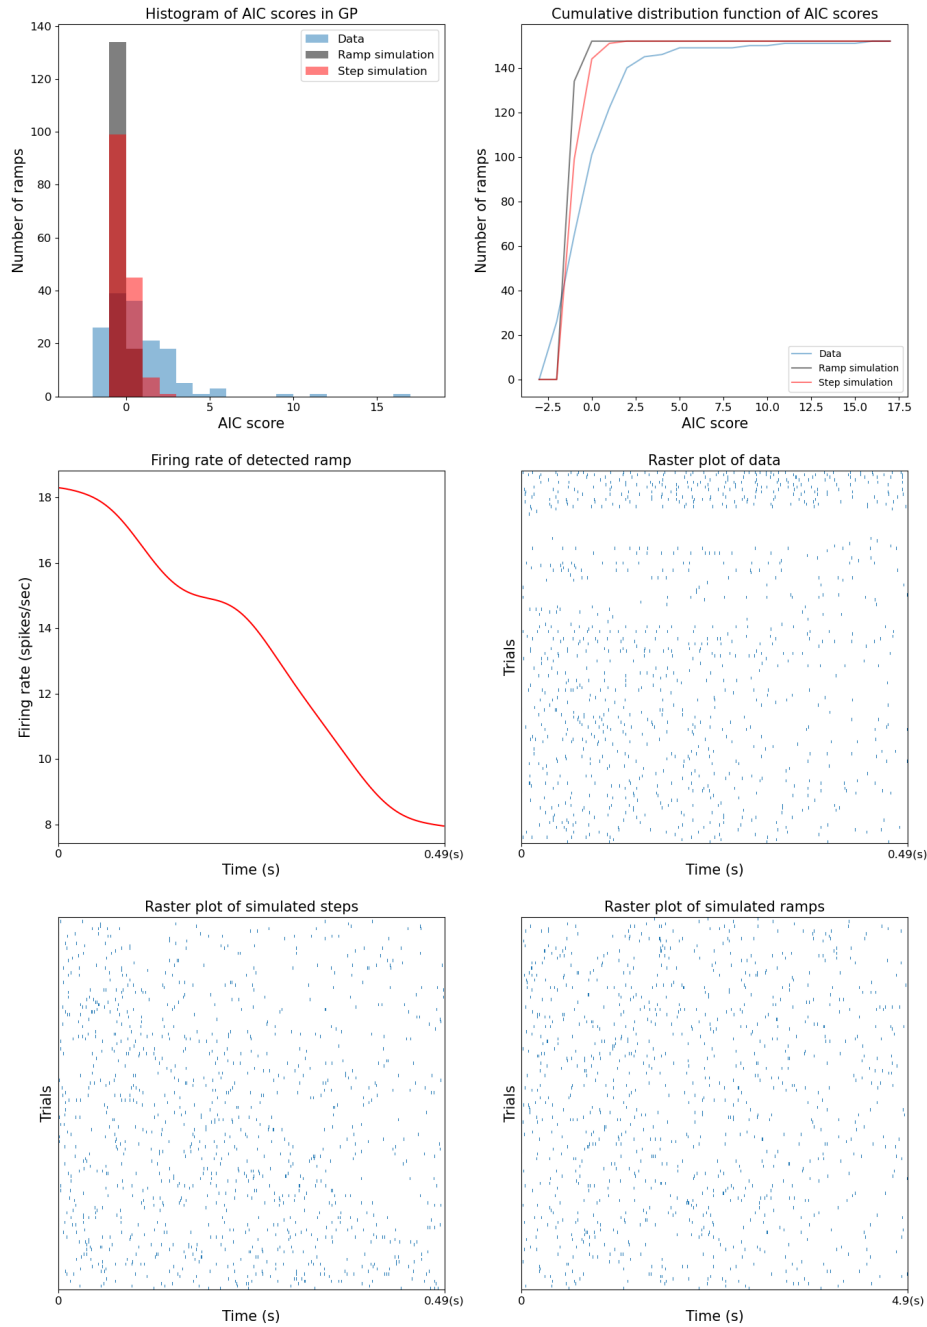


Figure B.27: Detected ramp in GP has ramping single-trial firing rates. Top row shows AIC score distribution on the left and the cumulative distribution of AIC scores on the right. Negative scores suggest the trial is more likely to be ramping while positive scores suggest the trial is more likely to be stepping. Middle row shows the firing rate of the detected ramp on the left with the raster plot of the trials on the right. Bottom row shows raster plots of simulated steps and ramps.

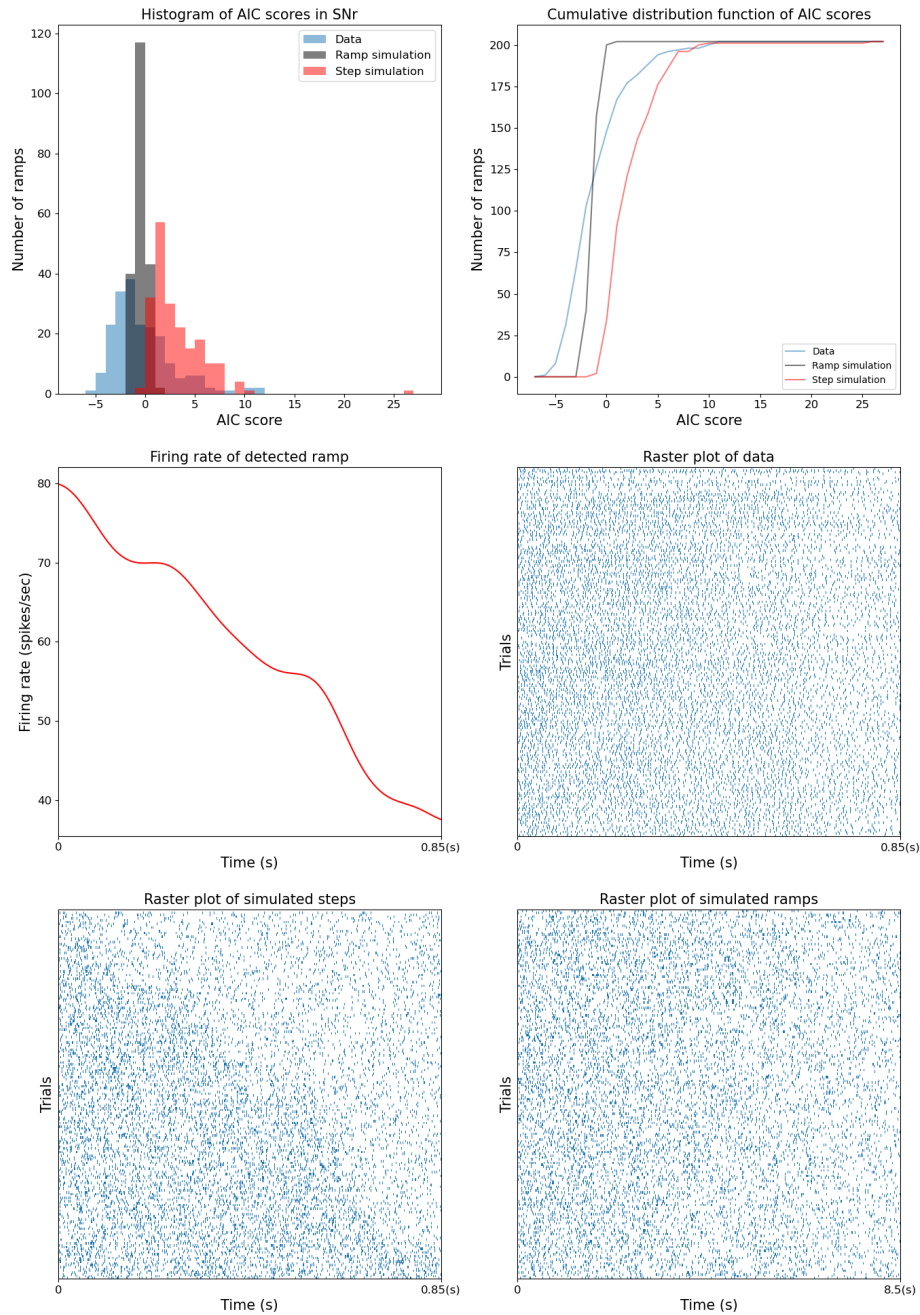


Figure B.28: Detected ramp in SNr has ramping single-trial firing rates. Top row shows AIC score distribution on the left and the cumulative distribution of AIC scores on the right. Negative scores suggest the trial is more likely to be ramping while positive scores suggest the trial is more likely to be stepping. Middle row shows the firing rate of the detected ramp on the left with the raster plot of the trials on the right. Bottom row shows raster plots of simulated steps and ramps.

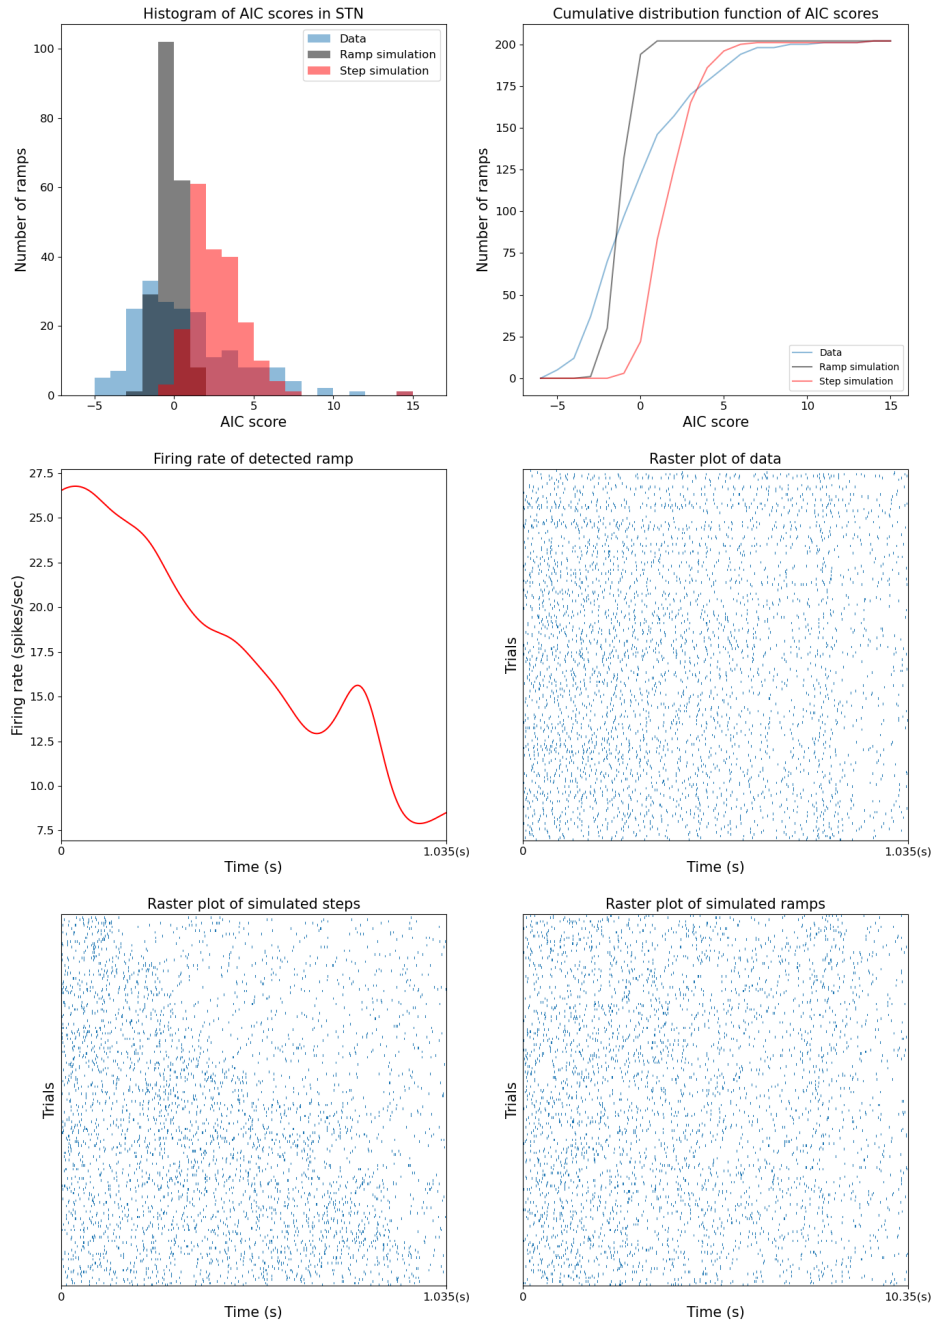


Figure B.29: Detected ramp in STN has ramping single-trial firing rates. Top row shows AIC score distribution on the left and the cumulative distribution of AIC scores on the right. Negative scores suggest the trial is more likely to be ramping while positive scores suggest the trial is more likely to be stepping. Middle row shows the firing rate of the detected ramp on the left with the raster plot of the trials on the right. Bottom row shows raster plots of simulated steps and ramps.

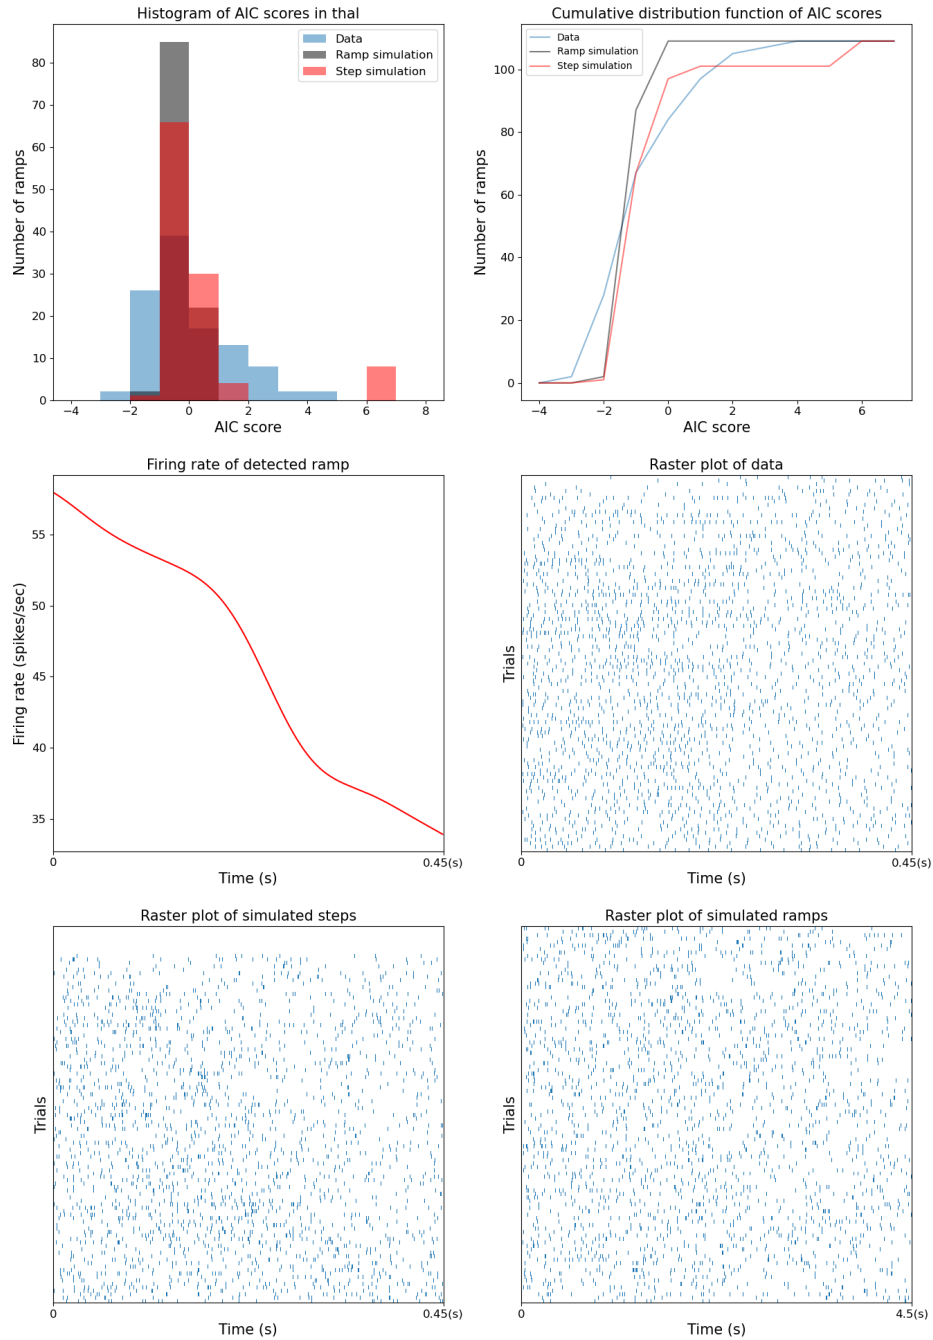


Figure B.30: Detected ramp in thal has ramping single-trial firing rates. Top row shows AIC score distribution on the left and the cumulative distribution of AIC scores on the right. Negative scores suggest the trial is more likely to be ramping while positive scores suggest the trial is more likely to be stepping. Middle row shows the firing rate of the detected ramp on the left with the raster plot of the trials on the right. Bottom row shows raster plots of simulated steps and ramps.