

Hit the brakes:

Unraveling oscillatory dynamics
during reactive stopping
in freely behaving rats



Jordi ter Horst

DONDERS
S E R I E S

**RADBOD
UNIVERSITY
PRESS**

Radboud
Dissertation
Series

Hit the brakes:

Unraveling oscillatory dynamics during
reactive stopping in freely behaving rats

Jordi ter Horst

Hit the brakes: Unraveling oscillatory dynamics during reactive stopping in freely behaving rats

Jordi ter Horst

Radboud Dissertation Series

ISSN: 2950-2772 (Online); 2950-2780 (Print)

Published by RADBOUD UNIVERSITY PRESS
Postbus 9100, 6500 HA Nijmegen, The Netherlands
www.radbouduniversitypress.nl

Design: Proefschrift AIO | Annelies Lips
Cover: Proefschrift AIO | Guntra Laivacuma
Cover image: Created in BioRender. Ter Horst, J. 2025
<https://BioRender.com/u72o926>
Printing: DPN Rikken/Pumbo

ISBN: 9789465150642
DOI: 10.54195/9789465150642
Free download at: <https://doi.org/10.54195/9789465150642>

© 2025 Jordi ter Horst

**RADBOUD
UNIVERSITY
PRESS**

This is an Open Access book published under the terms of Creative Commons Attribution-Noncommercial-NoDerivatives International license (CC BY-NC-ND 4.0). This license allows reusers to copy and distribute the material in any medium or format in unadapted form only, for noncommercial purposes only, and only so long as attribution is given to the creator, see <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

Hit the brakes:

Unraveling oscillatory dynamics during
reactive stopping in freely behaving rats

Proefschrift ter verkrijging van de graad van doctor

aan de Radboud Universiteit Nijmegen

op gezag van de rector magnificus prof. dr. J.M. Sanders,

volgens besluit van het college voor promoties

in het openbaar te verdedigen op

maandag 7 april 2025

om 14.30 uur precies

door

Jordi ter Horst

Promotor

Dr. M.X. Cohen

Copromotor

Dr. B. Englitz

Manuscriptcommissie

Prof. dr. R.J.A. van Wezel

Prof. dr. B.U. Forstmann, *Universiteit van Amsterdam*

Prof. dr. N. Böhler, *Universiteit Gent, België*

Contents

Chapter 1	General introduction	7
Chapter 2	Stopping speed as state, not trait: Exploring within-animal varying stopping speeds in a multi-session stop-signal task	23
Chapter 3	Decreased beta power and OFC-STN phase synchronization during reactive stopping	57
Chapter 4	General discussion	89
Appendices	References	100
	Nederlandse samenvatting	106
	Research data management	115
	About the author	116
	Portfolio	117
	List of publications	119
	Donders Graduate School for Cognitive Neuroscience	120



Chapter 1

General introduction

Reactive stopping

Imagine you are running on a pedestrian path and you approach a crossroad with traffic lights. The traffic light is green, so your brain has anticipated crossing the street. A couple of meters before physically passing the lights, the traffic light suddenly turns red. Now, the brain needs to abruptly instruct the muscles to stop the running movements of your legs to prevent a potentially dangerous scenario. This situation is called *reactive stopping*, as the planned and initiated movement needs to be stopped following an environmental change, which in this case is the traffic light turning red. It happens to be more difficult to stop in time the closer you get to the crossroad. The muscular effort of stopping itself may not be more difficult, but the time constraint makes it less and less likely for it to be in time (Logan & Cowan, 1984; Verbruggen & Logan, 2009). In other words, the timing between the red traffic light and your proximity to the crossroad is a limiting factor for being able to stop in time, independent from your general ability to abruptly stop prepared and initiated movements. Obviously, the further away you are from the crossroad while the traffic light turns red, the more likely it is you are able to stop in time.

Although stopping seems to require little mental effort, there are quite a lot of processes involved. First of all, the brain needs to perceive the stimulus from the environment and associate it with its contextual meaning. A red light on the ceiling in a disco will not let you think you should stop dancing, while a red light in a traffic situation will immediately be associated with stopping your (anticipated) movements. After the brain has perceived and associated the stimulus with the desired stopping action, it needs to plan and prepare the action for execution. Contrary to proactive stopping, where stopping is consciously anticipated and prepared before it is absolutely necessary, reactive stopping requires a quick unprepared response to an unforeseen event. This suggests that reactive stopping is carried out by a very quick and effective brain mechanism that is specific enough to cease certain muscle activity while maintaining posture and balance, without sacrificing speed of execution.

How reactive stopping is typically studied

In cognitive psychology and neuroscience, behavior is often studied using simplistic computer tasks in lab environments, where participants are presented with a stimulus that requires a certain response. Such a stimulus-response task




Figure 1. Typical stop-signal task for humans. After shortly fixating at a central position on a screen, a go-signal is presented instructing the participant to respond at the corresponding arrow direction side with a button press (go trial). Occasionally, this go-signal is followed by a stop-signal after a variable stop-signal delay (SSD), requiring the participant to not respond to the initially presented go-signal (stop trial). Taken and adapted from Figure 1 in Verbruggen et al. (2019).




Figure 2. Simplified illustration of the independent race between the go and stop process. Stopping fails when the response to the go-signal is finished earlier than the stopping process ($go\ RT < SSD + SSRT$), while stopping succeeds when the response to the go-signal is finished later than the stopping process ($go\ RT > SSD + SSRT$). As can be seen, the stop-signal delay affects stopping ability, given that the stopping time (SSRT) is fixed. SSD = stop-signal delay; SSRT = stop-signal reaction time; go RT = go trial reaction time. Taken and adapted from Figure 2 in Verbruggen and Logan (2009).

is repeated many times, and these repetitions are called *trials*. Reactive stopping is also studied using many repetitions of trials, typically in the stop-signal task (Figure 1). Participants are simply required to respond to a go-signal by pressing a button in front of them as quickly as possible. This go-signal is usually a visual stimulus, but can be any type of stimulus as long as it is associated with making a certain measurable response. These trials are called *go trials*. On a minority of trials (usually 25%), the go-signal is followed by another stimulus, the stop-signal, which requires the participant to stop the planned and initiated response. These trials are called *stop trials*. Similar to the traffic light example, the timing between the go-signal and the stop-signal determines the difficulty; the longer the delay between the go- and stop-signal, the more difficult it is to cease the initiated response to the go-signal in time (Figure 2). Conversely, it gets easier the shorter the delay is, because the response to the go-signal is still in the early phase and the response hand is more distant to the go-signal response button. This delay between the go- and stop-signal is called the stop-signal delay (SSD). Ideally the SSD is not fixed, as having a predictable delay for every stop trial could cause participants to wait for the stop-signal, and therefore succeed in stopping while they are in fact not responding to the go-signal in the first place. Therefore, the SSD is usually changed from trial to trial. When the participant succeeds in stopping the response to the go-signal, the SSD for the next stop trial is increased with a certain amount of milliseconds to make it a bit more difficult to stop in time on the next stop trial. The SSD is decreased when the participant fails to stop the initiated response to the go-signal, to make it a bit easier to stop on the next stop trial.

As the proportion of stop trials (25%) is three times smaller than the proportion of go trials (75%), stop-signal occurrence is somewhat unexpected and participants are therefore even more discouraged to wait for a possible stop-signal, as in most trials the stop-signal is not presented. Next to discouraging a waiting strategy, the adaptive SSD also allows for estimating the time it takes to stop, which is something that cannot be observed overtly due to the lack of a button press when a response is successfully stopped. As the SSD adapts after each trial based on stopping performance, the participant will eventually stop in about 50% of the stop trials, as the SSD changes towards delays where the participant is equally likely to succeed and fail at stopping. Following the principles of the independent horse-race model (Logan & Cowan, 1984; Verbruggen & Logan, 2009), the time it takes to stop (stop-signal reaction time; SSRT) can be estimated by finding the n -th fastest reaction time of the go trial distribution that matches with the probability of erroneously responding

while a stop-signal was presented (failed stopping), which is the go trial reaction time that matches with the internal stopping response (Figure 3). For example, when a participant succeeded in stopping in 55% of the stop trials, the probability of erroneously responding on a stop trial was 0.45 ($1 - 0.55$). The go trial reaction time matching this probability in a session containing 200 go trials is the 90th fastest go trial reaction time. The SSRT can then be estimated by subtracting the average SSD from the 90th fastest go trial reaction time, as the stop process only starts after the stop-signal (Verbruggen et al., 2019).




Figure 3. The independent horse-race model. The time it takes to stop (stop-signal reaction time; SSRT) can be estimated based on the go trial RT distribution, the SSD, and the probability of failed stopping. SSD = stop-signal delay; SSRT = stop-signal reaction time; RT = reaction time. Taken and adapted from Figure 2 in Verbruggen and Logan (2009).

Reactive stopping in the broader context of movement control

Reactive stopping is one of multiple subtypes of movement control. Movement control can generally be subdivided in two: initializing and executing movement, and suppressing movement. There are two classical pathways involved in orchestrating these two types of movement control: the direct and indirect pathway (Albin et al., 1989; Parent & Hazrati, 1995). Pathways can be thought of as specific routes between brain structures that allow for information transfer across them to support a certain function, such as the capability to sense, move or remember. The direct and indirect pathway are both situated in the basal ganglia, a group of deep brain structures important for various functions, of which the coordination of movement is the most well-known function. These pathways turn

out to be more complex than the description that will follow, but for explanatory purposes it will be kept a bit more digestible. The main function of the direct pathway is to initialize and execute desired voluntary movement, whereas the main function of the indirect pathway is to suppress undesired movement. In summary, the direct pathway works as follows: the cerebral cortex starts by generating the intention to move by sending activating (excitatory) signals to the striatum. Subsequently, the striatum suppresses the internal globus pallidus and substantia nigra pars reticulata. As the internal globus pallidus and substantia nigra pars reticulata normally suppress the thalamus, this suppression is now reduced, resulting in the thalamus sending activating signals to the motor cortex. Consequently, movement is facilitated. Conversely, the indirect pathway works as follows: the cerebral cortex sends activating signals to the striatum. However, in this case, the striatum suppresses the external globus pallidus. As the external globus pallidus normally suppresses the subthalamic nucleus, this suppression is now reduced. As a consequence, the subthalamic nucleus sends more activating signals to the internal globus pallidus and substantia nigra pars reticulata, leading to increased activity in the internal globus pallidus and substantia nigra pars reticulata. As a result, the internal globus pallidus and substantia nigra pars reticulata increase their suppressive signals to the thalamus, which in turn reduces activating signals to the motor cortex, leading to the prevention of movement (see Figure 4 for a schematic representation).

As both pathways are facilitated through activation of the striatum by the cerebral cortex, there needs to be a mechanism that enhances the function of either one of the pathways, depending on whether movement is wanted or not. This is managed by the release of a specific chemical substance called dopamine. Dopamine, one of the many neurotransmitters that can be found in our brain, is synthesized in the substantia nigra pars compacta. This dopamine is released in the striatum, close to specific brain cells called medium spiny neurons. The direct pathway mainly contains medium spiny neurons with type 1 dopamine (D1) receptors, while the indirect pathway predominantly contains medium spiny neurons with type 2 dopamine (D2) receptors. Critically, when dopamine binds to these two different receptor types, they have contrasting consequences for the two different pathways. While D1 receptor activation facilitates direct pathway functioning, D2 receptor activation suppresses indirect pathway functioning. As a result, dopamine release at medium spiny neurons in the striatum activates the pathway that facilitates movement, and suppresses the pathway that prevents movement from happening, together leading to movement (Surmeier et al., 2007; Gerfen & Surmeier, 2011).




Figure 4. A schematic representation of the direct and indirect pathway. As proposed by (Albin et al., 1989). Glutamate (red arrows), gamma-aminobutyric acid (GABA, gray arrows) and dopamine (purple arrows) play distinct roles in these pathways. D1 = dopamine receptor type 1; D2 = dopamine receptor type 2; GPe = external globus pallidus; STN = subthalamic nucleus; SNc = substantia nigra pars compacta; GPi = internal globus pallidus; SNr = substantia nigra pars reticulata; GABA = gamma-aminobutyric acid. Taken and adapted from Kandel et al. (2021).

In the context of reactive stopping, it may be thought that the indirect pathway is a perfect candidate for the implementation of ceasing ongoing movement, as an activated indirect pathway has a suppressing effect on movement. However, it is not likely that reactive stopping is executed by the indirect pathway. First of all, stopping ongoing movement is substantially different from suppressing potential new movements during immobility. In the first case muscles need to be instructed to do something else to prevent continuation of the ongoing movement – which may imply recruitment of additional muscles – while in the other case it is simply a matter of maintaining immobility. In addition, the

indirect pathway is a relatively long chain of brain structures connecting to each other (cerebral cortex - striatum - external globus pallidus - subthalamic nucleus - internal globus pallidus / substantia nigra pars reticulata - thalamus - motor cortex). Every time a signal needs to traverse from one brain structure to another, the travel time increases compared to when such a transition is not required. Thinking from an efficiency standpoint, this long chain of structures is not ideal for reactive stopping, as it usually needs to be executed very fast.

A more minimalist perspective suggests another possible candidate for the implementation of reactive stopping. You may only need one brain structure that is capable of recognizing the need to abruptly stop ongoing movements, which in turn signals another brain structure to implement the termination of ongoing movement. As excitatory thalamic projections to the motor cortex are associated with movement, it would make sense to suppress these signals during reactive stopping. As we know from the indirect pathway, a more activated subthalamic nucleus leads to a more activated internal globus pallidus and substantia nigra pars reticulata, which in turn suppress thalamic projections to the motor cortex. Thus, activating one of these three brain structures may be enough to terminate ongoing movement. As it turns out, more and more studies suggest that such a pathway, where one brain structure activates the subthalamic nucleus, internal globus pallidus or substantia nigra pars reticulata, may exist.

Evidence for a hyperdirect pathway

Many studies have shown that the subthalamic nucleus (STN) is important for stopping ongoing movement. For example, Aron and Poldrack (2006) used functional magnetic resonance imaging in humans to investigate which brain structures are activated during stopping in a stop-signal task. Among other parts of the brain, the STN was activated during stopping, and this STN activity was stronger for participants who had faster estimated SSRTs. Eagle et al. (2008) showed that lesions made in the STN of rats reduced the probability of stopping on stop trials, but go trial performance was unaffected. This effect was independent from the timing of the stop-signal delay, suggesting that a dysfunctional STN leads to a general stopping impairment. This was further supported by a study where brief optogenetic activation of the STN in mice caused abrupt pausing of licking behavior (Fife et al., 2017). Not only the STN is important for stopping, but Aron and Poldrack (2006) also showed that the




Figure 5. Brain regions related to reactive stopping. **A**, Functional magnetic resonance imaging (fMRI) showed that among other regions, the subthalamic nucleus (STN), pre-supplementary motor area (pre-SMA) and right inferior frontal cortex (rIFC) are activated during successful stopping. Taken and adapted from Aron and Poldrack (2006). **B**, Combining magnetoencephalography (MEG) and fMRI revealed that the pre-SMA and rIFC are involved in stopping. Taken and adapted from Schaum et al. (2021). **C**, Lesions in the ventral and lateral orbital cortex (VO/LO) led to slower stopping times (Eagle et al., 2008), and injections with norepinephrine agonists in the VO/LO increased stopping time speeds (Bari et al., 2011). Taken and adapted from Aron et al. (2014).

inferior frontal cortex, pre-supplementary motor area and globus pallidus display stop-related activity (Figure 5A). They found that participants with stronger stop-related inferior frontal cortex activity had stronger stop-related STN activity, and participants who quickly stopped their initiated actions did not only have stronger STN activity, as discussed before, but also had stronger stop-related inferior frontal cortex activity as compared to individuals that stopped at a slower pace. In line with this, Eagle et al. (2008) described the effects of lesions in rats to the orbitofrontal cortex (OFC, thought to be functionally comparable to the human inferior frontal cortex (Eagle & Baunez, 2010; Eagle et al., 2008; Parent & Hazrati, 1995)) and infralimbic cortex, and showed that only OFC lesioning caused SSRT slowing (Figure 5C). In addition, injection with a selective norepinephrine reuptake inhibitor (increasing signaling) in the OFC decreased stopping times without affecting going times or stop accuracy, as compared to a control injection (Bari et al., 2011). As the pre-supplementary motor area also seems to engage during stopping, researchers

questioned whether the inferior frontal cortex or pre-supplementary motor area is the first brain structure to take part in stopping. As functional magnetic resonance imaging does not provide sufficient temporal precision, Schaum et al. (2021) used a combination of magnetoencephalography and functional magnetic resonance imaging to elucidate which of the two is the initiator in stopping (Figure 5B). They revealed that inferior frontal cortex activity preceded pre-supplementary motor area activity related to stopping, and was correlated unidirectional from inferior frontal cortex to pre-supplementary motor area. Moreover, multiple tractography studies have shown that the inferior frontal cortex and STN connect through white matter (Aron et al., 2007; Neubert et al., 2010), and the integrity of white matter predicts stopping performance (Coxon et al., 2012; Forstmann et al., 2012).

But what does it mean when the brain is active?

Although very valuable for localizing which brain areas are involved, detected activation of brain structures does not tell much about how behavior and cognition are mechanistically implemented (Singh, 2012). And, activation patterns in functional magnetic resonance imaging studies are difficult to associate with particular events of the stop-signal task, as the hemodynamic response picked up with functional magnetic resonance imaging has temporal precision in the range of multiple seconds. This makes it hard to elucidate whether the measured increase in activity belongs to the presumably very quick and time-specific stopping behavior, or to something else related to or prior to stopping. Luckily, due to methods like electroencephalography, magnetoencephalography and electrocorticography, neuroscientists are able to look at brain dynamics with millisecond precision by measuring electrical (electroencephalography, outside the skull; electrocorticography, between brain surface and skull) or magnetic (magnetoencephalography, outside the skull) changes on the outside of the brain. Individual neurons generate very small electrical potentials and can not be detected with electroencephalography, magnetoencephalography or electrocorticography sensors due to the weakness of the potential. However, when large groups of neighboring neurons rhythmically synchronize their electrical activity, they together change the local electrical field as a whole, which can be observed with electroencephalography, magnetoencephalography or electrocorticography at the surface of the brain and skull. And, as the neighboring neurons activate together, they are also in their inactivated state together afterwards. As a

consequence, the local electrical field oscillates quickly between activated and inactivated states, better known as so-called brain waves. However, the brain contains many groups of neurons that all have their own individual brain waves, as different parts of the brain work in parallel to give rise to all the complex behaviors and cognitions we as humans are capable of. Some brain waves oscillate quickly, some slowly. And like waves in the ocean, brain waves can be big or small. All these individual processes happen simultaneously, and are picked up together at the sensors of the acquisition system. As a consequence, sensors receive a mix of many different brain wave sources at once. When sensors receive many different sources of brain activity, the signal is typically noisy and even uninterpretable without some clever processing and filtering.

Neuronal oscillations are associated with behavior and cognition

Brain waves, or neuronal oscillations, are observed during many different aspects of behavior and cognition. For example, alpha oscillations (8-12 Hz) are observed at the primary visual cortex during wakeful relaxation, especially when eyes are closed. For conscious perception, gamma oscillations (30-100 Hz) are thought to be important. Particular frequencies are not necessarily bound to particular parts of cognition or parts of the brain, as beta oscillations (12-30 Hz) are for example found in many different parts of the brain that support different cognitive functions. When we look at neuronal oscillations related to stopping, increased beta oscillations in the IFC and STN are hypothesized to be related to stopping (for an overview, see Aron et al. (2016)). For example, Swann et al. (2009) acquired electrocorticography data from four epilepsy patients and reported more beta power (13-18 Hz) at electrodes near the IFC in successful stop trials as compared to failed stop trials. Similarly, electroencephalography data revealed stronger right-frontal power in successful stop trials than failed stop trials in two of three studies, timed after the stop-signal but before the estimated SSRT (Wagner et al., 2018). In a group of medicated Parkinson's disease patients, Wessel et al. (2016) found that beta power in the STN was relatively stronger during successful stopping as compared to failed stopping, although this relative effect was only significant after SSRT. In line with this, deep brain stimulation in the STN of Parkinson's disease patients increased beta power at frontal scalp electrodes after the time of stopping and stopping speed increased as compared to when stimulation was off (Swann et al., 2011).

When neuronal oscillations are disrupted, they may affect normal functioning behavior and cognition, as excessive beta power in Parkinson's disease is associated with slowed movement (bradykinesia) and stiffness (rigidity). When beta power is reduced with treatments like levodopa and deep brain stimulation in the STN, symptoms can be alleviated in humans (Thomsen et al., 2020; Malvea et al., 2022). Parkinsonian motor behavior in rat models of Parkinson's disease can also be improved by optogenetic deep brain stimulation (Yoon et al., 2014; Yoon et al., 2016; Yu et al., 2020). All in all, these observations support the notion that beta activity is important for reactive stopping. Taken together with the previously mentioned functional magnetic resonance imaging and lesion studies, this may suggest that the IFC is the initiator of stopping by triggering the STN through beta synchronization, while the STN implements stopping (Aron et al., 2016), presumably by suppressing thalamic projections to the motor cortex through activation of the internal globus pallidus and substantia nigra pars reticulata.

What is synchronization and how could it facilitate communication across the brain?

Anatomical connectivity between brain areas does not change, at least not as fast as cognitive demands change from second to second. So, if the anatomical connectivity does not change so quickly, it must be something else that is quick and adaptive to support fast-changing cognition. Neuronal oscillations are quick as they can change at a millisecond timescale. But how could you interpret neuronal oscillations? As described before, when groups of neighboring neurons synchronously activate, the local electrical field starts to exhibit measurable electrical potentials at the surface of the brain or skull. The idea is that an oscillating group of neurons is in an 'active' state at the peak of an oscillatory brain wave, and in a more 'deactivated' or less active state at the trough of an oscillatory brain wave. Neuroscientists refer to the peaks of an oscillatory brain wave as windows of excitability, as this is the moment when the neurons are most susceptible to incoming electrical activity from other neurons, as well as more likely to send electrical signals to other neurons. Therefore, you could see neuronal oscillations as recurring rhythmic moments of possible communication (Fries, 2005). Synchrony between neurons does not only facilitate communication at a local scale. When two distant groups of neurons both exhibit oscillatory synchronization locally, they can also send and receive electrical signals from and to each other as two

separate groups. Critically, the two distant groups of neurons need to have an anatomical connection, and the time lag between the peaks of excitability of the two distant groups needs to line up with the distance and the speed with which the electrical activity travels. When these criteria are met, a sending group of neurons can effectively send a package of information (in the form of electrical signals) to a receiving group of neurons, as the information perfectly arrives when the receiving group of neurons' electrical activity is at peak excitability (Figure 6). In other words, long-range neuronal communication could be achieved through neuronal oscillatory coherence, which is why this hypothesis is called the communication-through-coherence hypothesis (Fries, 2005; Fries, 2015).

What is not understood about reactive stopping?

Although several studies seem to suggest that increased beta power is key to reactive stopping, there is not an abundant amount of electrophysiological evidence showing that both the IFC and STN oscillate in the beta range, specifically right after the instruction to stop an ongoing movement, and before stopping is executed. Some studies showed beta modulations after stop execution (Swann et al., 2009; Swann et al., 2011; Wessel et al., 2016; Hubbard & Sahakyan, 2023), and some before (Swann et al., 2012; Wagner et al., 2018). In addition, some studies either did not record from both areas at the same time, or lacked anatomical or temporal precision. Some studies had a low sample size, or effects seemed to be driven by a selection of participants. And, even when these two areas both oscillate at the beta frequency during stopping, it does not mean that they are communicating through coherence in the beta band, nor does it imply a causal relationship between beta band activity and reactive stopping. In this dissertation we adapted an existing rodent version of the human stop-signal task from Feola et al. (2000) and Bryden and Roesch (2015), and substantially improved some critical elements following recommendations from a consensus guide for the stop-signal task (Verbruggen et al., 2019), to make it as comparable to human stopping experiments as possible, and to improve the overall experimental quality of the rodent version of the stop-signal task. We trained male wild-type rats on this improved version of the stop-signal task, implanted them with custom-designed and custom-made electrodes in both the OFC and STN to record local electrical activity while they executed the stop-signal task. Not only does this allow for anatomical and temporal precision by recording local electrical

activity in the OFC (which is functionally comparable to the human IFC (Parent & Hazrati, 1995; Eagle et al., 2008; Eagle & Baunez, 2010)) and STN during stop-signal task execution; an animal model also allows future studies to causally interfere with methods like optogenetics to reveal whether the proposed hyperdirect pathway between the OFC and STN is causally related to reactive stopping.




Figure 6. Illustration representing the communication-through-coherence hypothesis. All green and yellow lines represent the excitability of individual groups of neurons in the visual cortex. The top and bottom line belong to groups of neurons in lower-order visual cortex, while the middle line belongs to a group of neurons in higher-order visual cortex. The top green line belongs to a group of neurons that represents an apple, while the bottom yellow line belongs to a group of neurons that represents a pear. Only information from the apple-representing group can be successfully transmitted to the higher-order visual processing group of neurons, as these packages of information arrive at peak excitability, contrary to the pear-representing group of neurons. Taken and adapted from Fries (2015).

Aim and outline of this thesis

In this thesis I investigated reactive stopping using dual-area, multi-electrode recordings in male rats performing a stop-signal task. The overall aim was to acquire more insight in the mechanisms underlying reactive stopping, and in particular within and between the orbitofrontal cortex and subthalamic nucleus, specifically after the instruction to stop and before stop execution. As we acquired many sessions of data per animal, we also investigated whether stopping speeds meaningfully changed from session to session. As such, in **chapter 2** we used multi-session stop-signal task data to investigate whether stopping speed is a fixed trait or a state that meaningfully varies from time to time. In addition, we assessed which factors played a role in single-session SSRT estimate reliability and which factors may explain changing stopping speeds. Equipped with the knowledge from chapter 2 that within-animal stopping speeds turned out to vary meaningfully from session to session, we investigated the role of neuronal oscillations in reactive stopping in **chapter 3** with high temporal and anatomical precision. Here, we recorded local electrophysiological activity in the OFC and STN during reactive stopping in the stop-signal task, as well as synchronization between the OFC and STN by using mathematical approaches that allowed for extracting long-range coherence in neuronal oscillatory activity. In **chapter 4** I summarize the key results of the thesis and discuss how they relate to existing literature on reactive stopping, point out the limitations of our studies and address future directions.

Chapter 2

Stopping speed as state, not trait: Exploring within-animal varying stopping speeds in a multi-session stop-signal task

Published as preprint:

Ter Horst, J., Cohen, M. X., & Englitz, B. (2024). Stopping speed as state, not trait: Exploring within-animal varying stopping speeds in a multi-session stop-signal task. *bioRxiv*. doi:10.1101/2024.09.05.611370

Abstract

Being able to reactively stop ongoing movements is important for safe navigation through the environment. Reactive stopping is typically studied using the stop-signal task, where participants are occasionally instructed to stop initiated movements. The speed of stopping, also referred to as the stop-signal reaction time (SSRT), is not observable because successful stopping lacks a response, but can be estimated. Researchers most often acquire one session of data per participant to estimate the speed of stopping, but sometimes more sessions of data are acquired to maximize the signal-to-noise ratio, for example when the task is combined with neural recordings such as electrophysiology. However, it is unknown whether the estimated stopping speed is a fixed trait or a state that can vary under identical experimental conditions. In this study, we investigate whether a separately estimated SSRT for each acquired session is statistically meaningful compared to estimating an across-session SSRT, by collecting many sessions in which male rats performed a stop-signal task. Results revealed that within-animal stopping speeds meaningfully changed from session to session and were not following a trend over time (e.g., due to task learning). Single-session SSRT estimates with lower reliabilities were associated with higher go trial response time variabilities, lower skewness levels of the go trial response time distribution, and lower stop accuracies. We also explored which factors explained changing SSRTs, and showed that motivation, shared motor dynamics, and attention could play a role. In conclusion, we encourage researchers to treat SSRTs as state-like variables when collecting multi-session stop-signal task data, as our results have convincingly shown that stopping speeds are far from trait-like under identical experimental conditions. This session-by-session approach will help future research in which neural signatures of reactive stopping need to be extracted in a time-precise manner, because time-locking stop-related neural activity to session-specific SSRTs is expected to capture the signature more precisely as opposed to an across-session SSRT.



Introduction

Reactive stopping in the real world

Being able to suddenly stop movements is a vital skill for navigating safely through the environment. Car drivers, bikers and other traffic users may suddenly cross your way while you are running around the block. Potential threatening external stimuli like these are taken very seriously by the brain; somehow it manages to reactively stop the ongoing running in a split second to prevent collision, without you even thinking about it. Reactive stopping, as defined here, is extremely fast and driven by external sensory input. The field of neuroscience has put many efforts into understanding how reactive stopping is implemented in the mammalian brain, but has not been conclusive about the specific neural mechanisms underlying the quick capability of stopping ongoing actions.




Figure 1. Typical stop-signal task for humans. After shortly fixating at a central position on a screen, a go-signal is presented instructing the participant to respond at the corresponding arrow direction side with a button press (go trial). Occasionally, this go-signal is followed by a stop-signal after a variable stop-signal delay (SSD), requiring the participant to not respond to the initially presented go-signal (stop trial). Taken and adapted from Figure 1 in Verbruggen et al. (2019).

Reactive stopping in the lab

Reactive stopping is typically studied with the stop-signal task, in which participants are presented with a short-lasting go-signal, requiring them to respond to this visual or auditory stimulus as quickly as possible by pressing

a button (Logan and Cowan, 1984; Verbruggen & Logan, 2009; Verbruggen et al., 2019). Occasionally, the go-signal is followed by a stop-signal, instructing to cancel the ongoing motor plan of pressing the go-signal button (Figure 1). This task can be anywhere from easy to very difficult, depending on the delay between the go-signal and the stop-signal. This delay is called the stop-signal delay. Comparable to a traffic light turning red while you were about to cross the road, a shorter delay (i.e., earlier stop-signal presentation) makes it easier to still stop the ongoing plan or action, while a longer delay (i.e., later stop-signal presentation) makes it harder to stop in time. In an experimental lab setting, successfully stopping an ongoing action does not come with an observable button press because participants have to cancel the ongoing move towards the button. Therefore, researchers need to estimate the non-observable stopping time, also called stop-signal reaction time (SSRT). This is done by making use of observable information such as the time needed to press when only a go-signal is presented, and the probability of (erroneously) responding with a button press when a stop-signal was presented (Verbruggen et al., 2019).

Is stopping speed a state or trait?

Quite often, researchers ask participants to perform the stop-signal task once and they need many participants for their study to acquire enough statistical power, as researchers are for example interested in whether a clinical group is slower at stopping than a non-clinical group. While this approach may be well-suited for the goal of the study, the downside is that it assumes that stopping speed is a trait that does not change from time to time. While cognitive development in childhood and adolescence is known to be associated with increased performance on stop-signal tasks, resulting in faster stopping speeds (Williams et al., 1999; Curley et al., 2018; Madsen et al., 2020), less is known about whether stopping speed is a fixed characteristic or a variable state after full cognitive development in non-clinical circumstances. Thunberg et al. (2024) have recently shown that test-retest reliability of SSRT estimates are low, even though SSRT estimates had high reliability within a session as demonstrated with high split-half reliability. However, to our knowledge there is no report out there that rigorously investigated how (in)variable stopping speeds are under identical experimental conditions in cognitively developed participants, and which factors drive potential instabilities in stopping speeds.

Animal models to study stopping

As opposed to comparing stopping speeds of a clinical group with a non-clinical group, researchers often have to turn to animal models when they

are interested in the neural mechanisms of stopping, especially when high spatial and temporal resolution are required. In these cases, researchers do not require many animals because within-subject, multi-session experimental designs are easier to perform with animals while maintaining enough statistical power. As a consequence, while many sessions are acquired, it remains unclear how one should handle multiple within-subject SSRT estimates. Are they statistically meaningful, reliable and useful?

Our experiment and key findings

To this end, we trained male rats ($N = 6$) on a rodent version of the stop-signal task and investigated whether within-animal single-session SSRT estimates were statistically meaningful as compared to just estimating an SSRT across all acquired sessions as if stopping speed were a fixed trait of the animal. Overall, we show that stopping speed is not a fixed trait, but a state that changes from time to time. Within-animal single-session SSRTs varied substantially over the course of many sessions, for some animals more than others. As compared to the within-animal, across-session SSRTs, single-session SSRTs were much less reliable. In addition, we explain which factors may play a role in different degrees of single-session SSRT reliability, and which cognitive and neural mechanisms may underlie changing stopping speeds.

Materials and Methods

Animals

Six wild-type Long-Evans rats participated in this study, aged 9 weeks and weighing 250-320 grams at the start of behavioral training (Charles River Laboratories, Calco, Italy). Rats were housed pairwise in Makrolon type III cages (UNO B.V., Zevenaar, The Netherlands) with a reversed 12-hour day-night cycle in a temperature- and humidity-controlled room ($21 \pm 2^\circ\text{C}$, $60 \pm 15\%$). As soon as the rats weighed more than 350 grams, they were housed in Makrolon type IVS cages to provide more horizontal space. As soon as the rats were implanted with electrodes (for another study, see chapter 3) they were housed individually, and the low conventional cage lid was replaced by a high cage lid to prevent damage to the implant. Corn cob granules were used as cage bedding, and sizzle bedding and a cardboard shelter were provided as cage enrichment. The rats were put on a restricted water intake schedule as soon as they acclimatized in the research facility. Every Monday to Friday morning the rats could get water in the behavioral task

(~5-8 mL, depending on performance), and in the afternoon they could drink *ad libitum* water for 30 minutes from a bottle. During weekend days, the rats received 30 grams of hydrogel (ClearH2O Inc., Westbrook, Maine, USA) each day, to keep the daily intake of water as stable as possible. Food pellets were provided *ad libitum* at all times. Weight and health were monitored on a daily basis. All animal procedures were approved by the Animal Welfare Body of the Radboud University Nijmegen and the Animal Experiment Committee (CCD No. AVD10300 2016 482, Project No. 2015-0129), according to national and international laws, to protect welfare under experimental conditions.

Skinner box

After acclimation of two weeks in the research facility the rats started with the restricted water intake schedule and behavioral training. Training and testing took place in a custom-built Skinner box (inside dimensions: 25 × 27 × 25 cm), with one wall containing three nose-poke ports (bottom-left, bottom-center, bottom-right, see Figure 2). Each port had an infrared emitter and phototransistor (type L-53F3C and L-53P3C, peak 940 nm, Farnell B.V., Utrecht, The Netherlands) enabling continuous automatic detection of a nose-poke by the rat. In addition, the left and right port also contained green light-emitting diodes (type L-53SGD-5V, peak 565 nm, Farnell B.V., Utrecht, The Netherlands) for presenting visual stimuli, as well as a small silicone tube at each bottom of the port for providing 50 µL water drop rewards driven by solenoid pumps (The Lee Company, Westbrook, Connecticut, USA). Electronics needed for the task in the Skinner box were controlled by a computer with custom-written code in MATLAB (R2018b, The MathWorks Inc., Natick, MA, USA).

Training procedure stop-signal task

The rodent version of the stop-signal task was based on the behavioral tasks used by Feola et al. (2000) and Bryden et al. (2012), with optimizations taken from Verbruggen et al. (2019). The training procedure consisted of three main steps, namely 1) central nose-poke initiation, 2) unilateral cue discrimination, and 3) stop trial introduction. Rats were trained for maximally one hour or 200 trials each day, while having rest during weekend days. During the central nose-poke initiation phase, the rats had to learn to initiate a trial by poking their nose in the central port. On the first day of training, poking at the central port for the shortest time detectable was already enough to initiate a trial, causing the presentation of a light cue (go-signal) for 100 ms either on the left or right side with an immediate water drop reward provided at the corresponding side. As the reward was provided immediately after trial initiation the rat was

not required to make a correct response yet. Despite not needing a correct response yet, trials were never terminated before a response was recorded, causing the rats to learn that they always had to respond in either the left or right port before they could initiate a new trial. This response requirement was kept at all training phases and data acquisition. The idea of this phase was to let the rat learn to associate central nose-poking with a positive outcome. Each day, the time required to stay in the central port (i.e., trial initiation time) was increased by 100 ms if the rat managed to initiate at least 100 trials in the previous session. When rats would release their nose from the central port before the required trial initiation time, the trial was not initiated and had to be re-initiated completely. This means premature responses as described in Verbruggen et al. (2019) are not possible. As soon as rats could initiate at least 100 trials in a session with 1000 ms initiation time, they proceeded to the second phase of training.

In the second phase, the rats had to learn to respond correctly to the go-signal before getting a reward. In practice this means the rat initiated a trial, received a go-signal at either the left or right side, and would only get a reward when the rat poked his nose in the port where the light was presented. This allowed the rats to learn to associate the go-signal side with the reward side. When the response was correct, the reward was immediately provided at the response port. No reward was provided after responding incorrectly. During this phase of training, the minimal time between the response at the lateral port and initiating a new trial (inter-trial interval) was gradually increased from 0 to 3 s in steps of 500 ms each session until 3 s was reached, to allow for proper separation of trials and to prevent rushing. To prevent go-signal anticipation, a jitter was slowly added to the trial initiation time from 0 to ± 200 ms in steps of 50 ms every following session until a jitter of ± 200 ms was reached. For each trial, the jitter value was randomly selected from a uniform distribution of numbers ranging from -200 to 200 ms with steps of 10 ms. As soon as the rats reached response accuracy above 80% for at least five consecutive days, they moved to the third and final phase of training.

In the final phase, stop trials were slowly introduced in addition to the go-only trials. On the first and second session of this final training phase the rats received 10% stop trials, the next three sessions 20% stop trials, and from the sixth day onwards 25% stop trials. In stop trials, the go-signal was followed by a stop-signal after a variable stop-signal delay (SSD). The stop-signal was given by a light on the other side than the go-signal side, and stayed illuminated

until the rat made a response at either the left or right port (Figure 2). The SSD was determined separately for left and right stop trials through a staircase procedure, where it increased with 50 ms in case of a correct response, and decreased with 50 ms in case of an erroneous response. The SSD was not separately determined for two animals, and for two other animals only after ~75% of their sessions were collected. This led us to exclude those four animals for this study, as we wanted consistency in the way the SSD was titrated for the particular purpose of this study. The starting SSD for stop trials was determined by averaging the SSDs from all stop trials from the previous session, for left and right stop trials separately. This staircase procedure is standard in human studies, ensures the collection of a wide range of SSDs and helps with obtaining a reliable stop-signal reaction time (SSRT) estimate (Verbruggen et al., 2019). Having a separately determined SSD for left and right stop trials ensures that the difficulty of a stop trial is always comparable between left and right stop trials, as a possible response bias may lead to an imbalanced accuracy for left versus right stop trials when there is a shared SSD. Go-signal side and stop trial occurrence were randomized and balanced as such that in each set of eight trials three left go trials, three right go trials, one left stop trial, and one right stop trial were randomly shuffled. Response bias was continuously checked by computing the percentage of left and right responses in the last 20 trials. When one side fell below 35%, the next trial was replaced by a go trial to that side to discourage response bias. As soon as the accuracy on stop trials floated around 50% and go trial accuracy was above 80% for five consecutive days, the rats were ready for electrode implantation. Training took approximately 6-8 weeks (30-40 sessions), and all rats achieved these criteria and were included in the subsequent recordings.

Data acquisition

During the workweek rats performed the stop-signal task daily for one hour or maximally 200 trials, while intracranial local field potentials were recorded from the OFC and STN. This electrophysiological data is part of chapter 3 and is therefore not further discussed here. Parameters such as trial initiation time, release time, movement time, and stop-signal delay (in case of a stop trial) were automatically saved to disk on the computer controlling the Skinner box. Only data acquired after electrode implantation was used for behavioral analyses, so training data is not incorporated.




Figure 2. Optimized rodent version of stop-signal task in a Skinner box. Rats initiated trials by poking in a central nose-poke port for 1000 ± 200 ms, whereafter they were presented with a short visual go-signal on either the left or right side. In 75% of trials, no other signal followed and the rat had to respond at the go-signal side for a water drop reward. However, in 25% of the trials, the go-signal was followed by a visual stop-signal after a variable delay (stop-signal delay, SSD), which instructed the rat to stop moving to the go-signal side and respond at the stop-signal side to obtain a water drop reward. Only correct responses were rewarded during data acquisition. RT = release time; MT = movement time; RespT = response time.

Behavioral data preparation

Analyses of behavioral data were done with custom-written code in MATLAB software (R2018b, The MathWorks Inc., Natick, MA, USA). For each trial, release time was defined as the time between go-signal onset and release from the central port, movement time was defined as the time between release from the central port and response at either the left or right port, and summed together they form response time. Importantly, response time is the homologue of reaction time in human stop-signal tasks, as in human tasks reaction time is defined as the time between go-signal and response. In addition, for stop trials the stop-signal delay (SSD) was defined as the delay between go-signal onset and stop-signal onset (Figure 2). To deal with anticipatory releases (although discouraged with a jitter for go-signal onset), a local minimum in the smoothed bimodal release time distribution was identified and used as the lower bound for trial removal, as this portion of releases was anticipatory instead of reactive and attentive to the go-signal. The maximal lower bound was set at 100 ms, and the upper bound was set at three standard deviations above the

mean release time. As no time limit was set for movement times, trials with movement times above five seconds were removed initially, whereafter the upper bound was set at three standard deviations above the mean for the remaining trials. Across the six included animals, trial removal based on release times and movement times resulted in an average removal of 15.1% of trials (95% CI [12.6 17.6]). The remaining trials were used for computing average release times, movement times, response times, stop-signal delays and stop-signal reaction times per session.

Session exclusion criteria

Sessions were excluded based on five different exclusion criteria. Sessions were excluded when: 1) the average response time on unsuccessful stop trials was larger than the average response time on go trials, as those sessions violated the independence assumption from the horse-race model (Verbruggen et al., 2019) – 3.3% (95% CI [1.0 5.6]); 2) the response accuracy on stop trials was lower than 25% or higher than 75%, for left and right stop trials separately – 8.9% (95% CI [5.0 12.8]); 3) there were less than 10 stop trials presented on each side – 10.2% (95% CI [4.9 15.5]); 4) the response accuracy on go trials is higher than the accuracy on stop trials, left and right trials separately – 6.2% (95% CI [4.2 8.2]); 5) the left or right SSRT estimate of that session was lower than 100 ms – 1.2% (95% CI [-0.3 2.7]). Across animals, these criteria led to removal of 22.5% of sessions (95% CI [17.6 27.4]).

SSRT estimation

Following the principles of the independent horse-race model, the unobservable stop-signal reaction time was estimated with the integration method (Verbruggen et al., 2019) for each session separately. In short, response times on go trials were sorted from shortest to longest, and the n -th response time matching the probability of failed stopping, $-p(\text{failed stopping})$, was found by multiplying $p(\text{failed stopping})$ from that session with the number of go trials in that session (see Figure 3 for a visual representation of go trial response times and how they relate to $p(\text{failed stopping})$, SSD and SSRT). Next, the SSRT was computed by subtracting the average SSD of that session from the n -th response time (nthRespTgo). As the SSD was separately determined for left and right stop trials, the SSRT was also separately estimated by using left and right response time distributions and SSDs. The response time distribution did not contain replacements for go omissions (go trials without a response), as responses were required on all trials. In addition, premature go responses (in our task releases before go-signal onset) are not part of the response

time distribution either, as trials are not initialized in those cases (see training procedure). Both correctly and incorrectly performed go trials were used for the distribution. The across-session SSRT was estimated by putting data from all sessions together as if it was one big session, and then handled like in the above-mentioned procedure. Confidence intervals for SSRT estimates were computed with bootstrapping, by selecting the n -th response time matching $p(\text{failed stopping})$ from that session from a randomly sampled but sorted response time distribution retrieved from go trials. This n -th response time was in turn used to compute the associated SSRT by subtracting the SSD of that session from the n -th response time. After 1000 bootstraps, the 2.5% and 97.5% percentile of this set of SSRTs were extracted for the lower- and upperbound of the 95% confidence interval. The same bootstrapping procedure is used for the 95% confidence intervals of the across-session SSRT estimates.




Figure 3. The independent horse-race model. The time it takes to stop (stop-signal reaction time; SSRT) can be estimated based on the go trial response time distribution, the SSD, and the probability of failed stopping. SSD = stop-signal delay; SSRT = stop-signal reaction time; RespT = response time. Importantly, RespT in our task is the homolog of reaction time in human stop-signal tasks, as in human tasks reaction time is defined as the time between go-signal and response. Taken and adapted from Figure 2 in Verbruggen and Logan (2009).

Trend analysis of SSRT estimates

As SSRT estimations on a single-session basis could imply trends over time, we utilized a Matlab function called *RobustDetrend* (version 9.4.0) to get a general idea of potential SSRT development during the course of multiple sessions, for each animal and stop-signal side separately. This function is able to derive the best polynomial fit in a series of data while preserving peak features (Schivre, 2024). The polynomial order limit was set at 10, but only

degree 0, 1 and 2 polynomials were found. We intentionally wanted to allow for peak features because SSRT estimations from sessions with smaller sample sizes may inherently have the tendency to show more legitimate outlier-like features, and we did not want to over-fit the data.




Figure 4. Fictive data to explain the different measures of statistical meaningfulness of single-session SSRT estimates. For each individual session, we checked whether 1) the single-session SSRT confidence interval (CI) was overlapping with the across-session SSRT estimate; 2) the single-session SSRT was within other single-session CIs; 3) the single-session SSRT was within neighboring single-session CIs. The purple thick line indicates the across-session SSRT with 95% confidence intervals indicated with thinner purple lines. Blue dots indicate single-session SSRT estimates with their corresponding 95% confidence intervals. Green arrows point to examples where the criteria are met as indicated, while red arrows point to examples where this is not the case. In this fictive example, criterion 1 would get a score of 60%, as three out of five sessions' confidence intervals were overlapping with the across-session SSRT (session 1, 3 and 5). Criterion 2 would reach a score of 45%, as four sessions had an SSRT positioned within two other confidence intervals (session 1, 2, 3 and 5) while one session (session 4) had an SSRT positioned within one of the other four confidence intervals ($4 \times 50\%$, $1 \times 25\% \rightarrow 45\%$). Criterion 3 would receive a score of 50%, because session 1 and 2 were within their neighboring confidence intervals (score of 100%), session 3 and 5 were not within neighboring confidence intervals (score of 0%), and session 4 was only within one neighboring interval (score of 50%), leading to an average of 50% ($2 \times 100\%$, $2 \times 0\%$, $1 \times 50\% \rightarrow 50\%$).

Statistical meaningfulness of single-session SSRT estimates

To address whether single-session SSRTs (or state SSRTs) were statistically meaningful or just numerically different from across-session SSRT estimates (or trait SSRTs) due to noise and limited sample size, we computed the

proportion of sessions where the single-session 95% confidence interval was overlapping with the across-session SSRT. Next, we investigated the degree of (dis)similarity between single-session SSRTs by computing the proportion of single-session SSRTs that were within the 95% confidence interval boundaries of other single-session SSRTs. As such, each session received a score between 0 and 100% and was used to compute an average across all single sessions. As a final step, the last analysis was repeated, but for neighboring sessions only, as this quantifies the speed at which proximate sessions can have significantly different SSRT estimates. For this, each session received a score depending on whether the single-session SSRT was not within neighboring 95% confidence intervals (score = 0%), within only one of them (score = 50%) or within both (score = 100%). In case of only one neighboring session (the first and last session), a score of either 0% or 100% was given (see Figure 4 for a visual representation of the three different analyses). All single-session percentages were averaged to obtain the percentage as illustrated in Figure 7.

These three different proportion calculations were repeated with detrended SSRT data to analyze if and how much proportions were affected by trends in SSRTs. We refrained from statistically testing whether detrending SSRTs affected these proportions because tests were heavily underpowered due to low sample size ($N = 6$). We qualitatively compared the proportions derived from unfiltered and detrended SSRTs.

Reliability of single-session SSRTs

How certain one can be about a single-session SSRT estimate is quantified by the 95% confidence interval. Although it is inevitable to have larger confidence intervals as compared to the confidence obtained with across-session SSRTs, it is still worth investigating which factors were associated with smaller windows of confidence and hence more reliable single-session SSRT estimates. For each animal and trial side (left/right) we fitted a separate linear regression model with session-specific, independent variables: variability of go trial response times, skewness of the go trial response time distribution, stop-signal delay stabilization, stop accuracy and number of stop trials, and the 95% confidence interval width of SSRT as dependent variable. The model was fitted with the built-in Matlab function *fitlm* and included the intercept and main effect terms. As ordinary least squares models are highly sensitive to outliers, the model performed robust regression using the bisquare weighting function. This made the model less susceptible to disproportionate leverage from outliers. Variability of go trial response times was defined as the root mean

square of demeaned go trial response times. Skewness of the go trial response time distribution was found with the built-in Matlab function *skewness*. Stop-signal delay stabilization was defined as the root mean square of demeaned stop-signal delays in the second half of the session. Stop accuracy was defined as the percentage of correct responses on stop trials, and lastly, the number of stop trials as the amount of stop trials that were originally in the uncleaned dataset. All variables were normalized to z-scores, so we could obtain standardized beta coefficients that were directly comparable.

To be able to visualize the relationship between each individual independent variable with the SSRT 95% CI width, we extracted the intercept and slope from simplified robust linear regression fits as described before, but with only one independent variable in the model. As such, we could include a linear line for each pair of independent variable and dependent variable in Figure 8A, for each animal separately. Lines were only solid when the independent variable in the full robust linear regression model had a p-value less than .05, and was dashed in case of non-significance. Standardized beta coefficients from all five independent variables and all six animals were represented with colors and summarized in heatmaps for left and right trials separately. A two-sided Wilcoxon signed-rank test was performed for each independent variable separately, and tested whether the standardized beta coefficients were significantly different from 0.

SSD and nthRespTgo contribution to SSRT estimates

Because the non-observable SSRT was computed by subtracting the SSD from nthRespTgo, similar SSRTs could be obtained with many different combinations of SSD and nthRespTgo. Likewise, different SSRTs could be obtained with a constant SSD or nthRespTgo, while the other variable changes from session to session. To better understand the relationship between SSD, nthRespTgo and SSRT, we visualized these three variables together in Figure 9. We plotted the observable variables on the x- and y-axis, while using a diagonal for the across-session SSRT to illustrate how different sessions (and their corresponding SSD and nthRespTgo) relate to each other and the across-session SSRT. The average difference between single-session SSRTs and the across-session SSRT was computed for each animal and trial side (left/right) as a proxy for how spread out the single-session SSRT estimates were.

To statistically quantify the previously mentioned relationship, and to find out how much the SSD and nthRespTgo each contributed to session-to-session SSRT variance, we fitted a linear regression model for each animal and trial

side (left/right) separately with SSRT as dependent variable and SSD and nthRespTgo as independent variables. The model was fitted with the built-in Matlab function *fitlm*, included the intercept and main effect terms, and made use of the bisquare weighting function to make the model more robust to outlier sessions. Standardized beta coefficients were extracted by z-scoring all input variables, so we could directly compare how much each variable's variance contributed to session-to-session SSRT changes. A paired-samples t-test was used to statistically test whether there was a difference in how much each of two independent variables affected the SSRT. To this end, the magnitudes of standardized beta coefficients (i.e., ignoring the sign of the coefficients) from all animals' robust linear regression models were used as paired samples.

Possible cognitive and neural mechanisms driving session-to-session SSRT variability

Next to the obvious numerical predictors for session-to-session SSRT changes, SSD and nthRespTgo, we thought about which cognitive and neural mechanisms could possibly cause animals to have changing stopping speeds from session to session. Because the stop-signal task was optimized for estimating the SSRT and extracting neural stopping signatures, and not for investigating which factors contribute to within-animal SSRT variability across sessions, we approximated three different cognitive and neural mechanisms for changing SSRTs with the data that were available from the task: 1) motivation for quick rewards; 2) shared motor dynamics; 3) attention. Other possible contributing factors will be discussed theoretically in the discussion.

When animals are motivated for a quick reward, one would expect that they release fast after seeing the go-signal. Therefore, motivation for quick rewards was approximated with the release time on go trials (RTgo), and tested with linear regression by means of the built-in Matlab function *fitlm*, with a bisquare weighting function. Since this motivation could be expected at any stop-signal delay, the sessions were divided into four equally sized groups based on the average session SSD: 0-25th SSD percentile, 25-50th SSD percentile, 50-75th SSD percentile, and 75-100th percentile. A separate model was fitted for each trial side and SSD percentile group, with SSRT as dependent variable and RTgo as independent variable. Variables were standardized to obtain comparable standardized beta coefficients.

If it were true that going and stopping share motor dynamics, it is reasonable to expect that sessions where animals needed more time to respond on go

trials would also have longer stopping times. For that reason, we tested the hypothesis of shared motor dynamics with linear regression for each trial side separately (*fitlm*, with bisquare weighting function), with *RespTgo* as independent variable and *SSRT* as dependent variable. Variables were standardized, so we could average the standardized beta coefficients among animals with significant regressions.

When animals are attentive they are likely to respond quickly to the initial signal and perform well as they are paying attention to the location of the stimulus. As such, a possible effect of attention was tested by comparing *SSRTs* between sessions with high average *RTgo* and low response accuracy on go trials (low attention) and sessions with low average *RTgo* and high response accuracy on go trials (high attention). Sessions belonged to the low attention group if the average *RTgo* was higher than the median across all sessions (median $RTgo_{\text{left}} = 254.0$ ms; median $RTgo_{\text{right}} = 234.9$ ms) and the average response accuracy on go trials was lower than the median response accuracy across all sessions (median $Ago_{\text{left}} = 86.0\%$; median $Ago_{\text{right}} = 87.6\%$). Sessions belonged to the high attention group if the average *RTgo* was lower than the median across all sessions and response accuracy on go trials was higher than the median response accuracy across all sessions. A t-test for random samples was used for each trial side separately to statistically test whether attention (low vs. high) had a significant effect on *SSRT*.

Results

In the following sections, we will address three main questions. 1) Are single-session *SSRTs* statistically meaningful as compared to the across-session *SSRT*? 2) Which factors play a role in the reliability of single-session *SSRT* estimates? 3) Which factors contribute to varying *SSRTs*?

Trends in single-session *SSRT* estimates

From visual inspection (Figure 5) it is clear that the single-session *SSRTs* varied over sessions (each session is a different testing day). We first asked whether the *SSRTs* followed a temporal trend that could account for this variability, e.g., if *SSRTs* decrease over time as the animals were in the task for longer (but note that the animals were fully trained before any of the data shown here). As single-session *SSRT* estimates seemed far from constant, we investigated whether *SSRT* estimates followed a trend as the number of




Figure 5. Single-session stop-signal reaction time (SSRT) estimates for each animal, with corresponding 95% confidence intervals. Blue-colored dots represent SSRTs belonging to left stop trials, while red-colored dots represent SSRTs belonging to right stop trials. Vertical blue and red thick lines are across-session SSRT estimates (in case all sessions within an animal would have belonged to one big session) with 95% confidence intervals indicated with thinner vertical lines. Each animal is labeled with a color that matches color-use in other figures.

acquired sessions progressed. Polynomial fitting revealed only non-zero constants, linear and quadratic trends, as additional polynomial orders were not statistically improving the fit (Figure 6). For some animals (blue, orange, brown) SSRTs followed a negative quadratic trend (n-shaped), while for other animals (green, brown) the SSRTs followed a positive quadratic trend (u-shaped). For one animal (brown) the sign of the quadratic fit was opposite between left and right SSRTs, while for another animal (green) the sign was similar. Left SSRTs were constant for one animal (yellow), and right SSRTs were constant for two animals (black and blue). After all, SSRT trends were not comparable between animals and none of them had similar kinds of polynomial fits for left and right SSRTs.




Figure 6. Single-session stop-signal reaction time (SSRT) estimates show different trends over time for different animals and stop-signal sides (**A**, left SSRTs; **B**, right SSRTs). Each dot represents a single session from one animal, where different colors indicate different animals. Trendlines, illustrated with thick colored lines, are polynomials with either a non-zero constant, linear or quadratic fit. For left SSRTs, data from animal #4 (yellow) was best explained with a constant SSRT, while for right SSRTs, data from animal #1 (blue) and #2 (black) were best explained with a constant SSRT. All other groups of SSRTs were better explained by either a linear decline, or negative and positive quadratic polynomials. Note that the polynomial fits are not all spanning the same amount of sessions, because the amount of sessions was not identical between animals.

Statistical meaningfulness of single-session SSRT estimates

Visual inspection of Figure 5 suggests that individual session SSRTs might be significantly different from the across-session SSRT. To statistically address this, we checked whether confidence intervals of single-session SSRTs overlapped with the across-session SSRT. Confidence intervals of single-session SSRTs overlapped with the across-session SSRT in 47.2% (95% CI [42.9 51.5]) and 43.1% (95% CI [35.4 50.8]) of the sessions for left and right, respectively (see unfiltered SSRT proportions in Figure 7 for individual animal proportions). This means that more than half of sessions had a significantly

different SSRT than the across-session SSRT. It is possible that the results described above were simply due to a trend over time, for example if SSRTs decrease over time as animals gained expertise with the task (but note that the animals were fully trained before we started collecting data). Therefore, we applied detrending to the SSRT data with the polynomial fits we showed before (Figure 6) so we were able to correct for trends in SSRTs. Detrending the SSRTs numerically increased the overlap between confidence intervals of single-session SSRTs with the across-session SSRT to 55.2% (95% CI [48.7 61.8]) and 51.5% (95% CI [46.2 56.8]) for left and right, respectively (see detrended SSRT proportions in Figure 7 for individual animal proportions). This and following comparisons between unfiltered and detrended SSRT proportions were not statistically tested because of low sample size ($N = 6$), which made tests heavily underpowered.

To find out how (dis)similar single-session SSRTs were, we computed the average proportion of sessions where the single-session SSRT was positioned within the 95% confidence interval of other single-session SSRT estimates. For unfiltered SSRT data, only 31.3% (95% CI [34.8 27.8]) of left and 28.8% (95% CI [24.4 33.2]) of right single-session SSRTs were statistically similar to other single-session SSRTs, indicating that the majority of single-session SSRTs were not coming from the same distribution. Detrending the SSRTs numerically increased the overlap of neighboring sessions to 36.7% (95% CI [32.0 41.4]) and 31.7% (95% CI [28.0 35.3]) for left and right, respectively.

As one might expect proximate or neighboring sessions to have comparable SSRTs when SSRTs are slowly changing over time, the latter analysis was also applied for neighboring sessions only to investigate the degree to which neighboring sessions are (dis)similar. With the unfiltered SSRT data, this resulted in an increased fraction of SSRTs positioned within single-session confidence intervals, i.e. 43.6% (95% CI [40.1 47.1]) of left and 38.3% (95% CI [33.9 42.7]) of right single-session SSRT estimates were statistically similar to other single-session SSRTs. Although this increase indicates that neighboring sessions had more comparable SSRT estimates than distant sessions, it still shows that more than half of the single-session SSRTs were statistically significantly distinct from the surrounding SSRTs. Detrending the SSRTs numerically increased this to 44.1% (95% CI [37.6 50.7]) for right SSRTs, and numerically decreased to 35.5% (95% CI [27.4 43.6]) for left SSRTs.




Figure 7. Proportion of sessions meeting three different criteria. **A**, Proportion of sessions where the single-session confidence interval (CI) is overlapping with the across-session stop-signal reaction time (SSRT). **B**, Proportion of sessions where the single-session SSRT is positioned within all other single-session CIs. **C**, Proportion of sessions where the single-session SSRT is positioned within neighboring single-session CIs (see Methods). Different colors represent different animals, and the stop-signal side is separated in two different columns (left/right). Within each left and right separation, two proportions are connected with a colored line. The left proportion is computed without any detrending (unfiltered SSRTs), while the right proportion is computed after detrending single-session SSRT estimates (detrended SSRTs).

Reliability of single-session SSRT estimates

Estimating the non-observable SSRT from smaller than ideally-sized sessions came with less reliable estimations. Separately computing the left and right SSRT reduced sample size even more for each SSRT estimation. Therefore, we explored which factors relate to the degree of reliability of the SSRT estimate, by using the width of the 95% confidence interval as a dependent variable in a set of linear regression models for each animal and trial side separately (Figure 8). Across both trial sides, three out of five variables showed a significant group effect: variability of go trial response times (mean $\beta^*_{\text{left}} = .47$, $p = .031$; mean $\beta^*_{\text{right}} = .53$, $p = .031$), skewness of go trial response time distribution (mean $\beta^*_{\text{left}} = -.35$, $p = .031$; mean $\beta^*_{\text{right}} = -.33$, $p = .031$) and stop accuracy (mean $\beta^*_{\text{left}} = -.33$, $p = .031$; mean $\beta^*_{\text{right}} = -.40$, $p = .031$). All animals had an individual significant effect of variability of go trial response times for left and right trials, while go trial response time distribution skewness was not




Figure 8. Robust linear regression with stop-signal reaction time (SSRT) 95% confidence interval (CI) width as a proxy for SSRT estimate reliability. **A**, Scatterplots for each independent variable (RespTgo variability, RespTgo skewness, SSD stabilization, stop accuracy, number of stop trials) and SSRT CI width, for left and right trials separately. Each dot represents a session and colors represent different animals. Linear lines are extracted from simplified robust linear regression models (see Methods), and are solid when $p < .05$ in the full robust linear regression model, and dashed in case of non-significance. **B**, Summary heatmaps for left and right trials separately, in which each of first six rows represents an animal, containing the standardized beta coefficients for each independent variable (columns) in the full robust linear regression model. Asterisks indicate $p < .05$ of that predictor variable in robust linear regression, and match with solid lines in panel A. Last row in each heatmap contains average standardized beta coefficients across all six animals, and asterisks indicate $p < .05$ in the Wilcoxon signed-rank test. Colored bars on the left of the heatmap represent different animals and match the color-scheme of panel A and other figures in this chapter. Adjusted R^2 values from the full model are shown on the right side of the heatmap. RespTgo = response time on go trials; SSD = stop-signal delay.

significant in three models: 1st animal left (dark blue), 1st animal right (dark blue) and 5th animal right (orange). Stop accuracy did not have an individual significant effect in one model: 3rd animal right (green). Despite the absence of these four individual significant effects, the signs of standardized beta coefficients were in line with the other animals. For all animals, at least two out of five, but not more than three out of five, variables significantly contributed to the explained variance of SSRT reliability.

The Wilcoxon signed-rank test for group significance was significant only when signs of standardized beta coefficients were consistent across all six animals, which resulted in an additional significant group effect for number of stop trials for right trials (mean $\beta^*_{\text{right}} = -.18, p = .031$). However, only one animal had individual significance for this variable. On average, roughly half of the SSRT's variance in reliability was explained by the robust linear regression model (adj. $R^2_{\text{left}} = .48$; adj. $R^2_{\text{right}} = .51$). On an individual animal basis, all robust linear regression models were significantly different from a constant model (all $p < .05$) except for the right trial model of the 5th animal (orange; $F_{\text{right}}(31, 25) = 2.24, p = .082$). In line with this, the lowest amount of SSRT reliability variance was explained for the 5th animal (orange; mean adj. $R^2 = .38$), while the highest amount of variance was explained for the 2nd animal (light blue; mean adj. $R^2 = .68$).

SSD and nthRespTgo contribution to SSRT estimates

Higher values of nthRespTgo were not necessarily accompanied with higher SSDs, and higher SSDs were not always accompanied with higher values of nthRespTgo (Figure 9). Thus, slower movement speeds were not associated with later stop-signal onsets and vice versa. This is exactly what could have been expected from session-to-session changing SSRTs. When SSRTs would have been similar across sessions, a changing movement speed would have always been perfectly compensated with a changing SSD to obtain a similar SSRT (resulting in dots lined-up on the diagonal across-session SSRT in Figure 9). The average difference across animals between single-session SSRTs and the across-session SSRT was 73.9 ms (95% CI [58.6 89.2]) and 71.3 ms (95% CI [50.4 92.1]) for left and right trials, respectively.

SSRT was computed by subtracting SSD from nthRespTgo. Since both observable variables had considerable spread across sessions, it was not a surprise that robust linear regression revealed that both variables were significant predictors for SSRT for all animals and trial sides (all $p < .001$). The main goal of fitting the robust linear regression model was to extract

standardized beta coefficients, so we could directly compare each variable's contribution to SSRT. While for some animals absolute standardized beta coefficients were numerically bigger for nthRespTgo than for SSD, for some animals this was the opposite (mean $\beta_{\text{nthRespTgo, left}}^* = 0.86$ (ranging from 0.80 to 0.95), mean $\beta_{\text{SSD, left}}^* = -0.79$ (ranging from -1.15 to -0.50); mean $\beta_{\text{nthRespTgo, right}}^* = 0.85$ (ranging from 0.59 to 0.94), mean $\beta_{\text{SSD, right}}^* = -0.83$ (ranging from -0.92 to -0.72)). A paired-samples t-test indicated that there were no statistically significant differences between absolute standardized beta coefficients of SSD and nthRespTgo, for both left and right trials ($t_{\text{left}}(5) = .87, p = .424$; $t_{\text{right}}(5) = .17, p = .873$), meaning SSD and nthRespTgo had statistically comparable contributions to SSRT estimates.




Figure 9. Stop-signal delay (SSD) and nthRespTgo in relation to the across-session stop-signal reaction time (SSRT) estimate. Each animal is indicated with a different color. Each dot is a session and corresponds to a single-session SSRT estimate ($\text{SSRT} = \text{nthRespTgo} - \text{SSD}$). X- and y-axis limits are identical within each animal, so dot positions and across-session SSRTs between trial sides (left/right) are directly comparable. The diagonal line represents the across-session SSRT estimate (value also printed close to each diagonal line), and includes 95% confidence intervals with thinner diagonal lines. Orthogonal distance from dot to diagonal SSRT line represents deviance from across-session SSRT. As such, all positions on the diagonal line are equal to the across-session SSRT, while dots left of and above the line represent sessions with a lower SSRT than the across-session SSRT, and dots right of and below the line represent sessions with a higher SSRT than the across-session SSRT. Width of the shaded area is equal to 100 ms (50 ms on each side of across-session SSRT). nthRespTgo = n-th response time on go trials matching p(failed stopping).

Possible cognitive and neural mechanisms driving session-to-session changes in SSRT

Because SSRTs changed from session to session, we reflected on which cognitive and neural factors could drive these changing SSRTs. We identified three, in our opinion, most likely contributors: Motivation for quick rewards, shared motor dynamics and attention. We attempted to quantify these cognitive and neural mechanisms with variables from the task, and theoretically addressed other possible driving factors in the discussion which we could not approximate with task variables.

Motivation for quick rewards

It is possible that sessions with faster SSRTs were those in which animals were more motivated to obtain rewards (e.g., they were thirstier in some sessions). If this were a viable explanation, it would predict that the release time on go trials (RTgo) would strongly correlate with SSRT. Therefore, we regressed RTgo with SSRT in four equally sized SSD quantiles across animals for left and right trials separately (Figure 10A). Almost all quantiles showed a significant positive association between RTgo and SSRT (except for 3rd and 4th quantile of left trials, all other $p < .05$). Across trial sides (left/right), standardized beta coefficients and adjusted R^2 s were smaller for each next SSD quantile ($\beta^*_{1st} = 0.61$, $\beta^*_{2nd} = 0.58$, $\beta^*_{3rd} = 0.32$, $\beta^*_{4th} = 0.16$; adj. $R^2_{1st} = 0.41$, adj. $R^2_{2nd} = 0.31$, adj. $R^2_{3rd} = .09$, adj. $R^2_{4th} = 0.04$), suggesting that high motivation for quick rewards (as reflected with low RTgo) was associated with fast stopping, but mostly in sessions where stop-signals were presented relatively early, and not as much when stop-signals were presented later.

Shared motor dynamics

As elements of motor dynamics between going and stopping could be shared, we assessed whether RespTgo and SSRT were positively associated (Figure 10B). For left trials, 4 out of 6 animals had a significant positive association between RespTgo and SSRT (mean $\beta^*_{left} = 0.63$, all $p < .05$), while for right trials only 2 out of 6 animals had a significant positive association (mean $\beta^*_{right} = 0.59$, all $p < .05$). Therefore, general motor dynamics were likely also a contributor to across session SSRT variability.




Figure 10. Motivation, motor dynamics and attention contribute to session-to-session changing SSRTs. **A**, Linear regression between RTGo and SSRT for 1st, 2nd, 3rd and 4th percentile group of session SSDs across all six animals (therefore no colored dots). Dots represent different sessions. Left and right trials separated by top and bottom row, respectively. Grey line represents the robust linear fit. Adjusted R^2 printed in right bottom corner if $p < .05$, otherwise non-significant (n.s.). **B**, Linear regression between RespTgo and SSRT for each animal and trial side separately. Dots represent different sessions. Different animals are indicated with different colors, with the same color scheme as the other figures in this chapter. Dashed lines indicate non-significant regressions, while solid lines indicate significance ($p < .05$). **C**, Across animals, sessions with high attention have significantly lower SSRTs as compared to sessions with low attention. Blue dots are left trials, red dots are right trials. Black horizontal lines indicate the average SSRT of that group. Asterisks indicate significance ($p < .01$). RTGo = release time on go trials; SSRT = stop-signal reaction time; RespTgo = response time on go trials.

Attention

We approximated the level of attention with a combination of two variables: RTgo and response accuracy on go trials (Figure 10C). Sessions in which animals had relatively low attention (high RTgo and low accuracy on go trials) had significantly higher SSRTs (mean SSRT_{left}: 371.1 ms; mean SSRT_{right}: 451.3 ms), as compared to sessions where animals had relatively high attention (low RTgo and high response accuracy; mean SSRT_{left}: 296.6 ms; mean SSRT_{right}: 319.7 ms) ($t_{\text{left}}(54) = 3.2583, p < .01$, Cohen's $d = 0.87$; $t_{\text{right}}(50) = 5.0951, p < .001$, Cohen's $d = 1.41$). Therefore, sessions with higher levels of attention were also associated with faster stopping speeds.

Discussion

In this study, we collected many sessions per animal in a rodent version of the stop-signal task, and explored whether single-session SSRT estimates were statistically meaningful as compared to a single SSRT estimate across all sessions. We showed that stopping speeds were changing from session to session and did not show a basic or consistent trend over time across individuals. As single-session SSRT estimates came with higher degrees of estimate uncertainty, we wanted to know which factors were underlying these increasing uncertainties. Higher degrees of single-session SSRT reliability were associated with lower go trial response time variabilities, lower skewnesses of the go trial response time distribution and higher stop accuracies. SSRT variability is equally explained by variability in stop-signal timing and go trial response times, and motivation, shared motor dynamics and attention could partly explain changing stopping speeds. This study suggests that stopping speed is a state that can meaningfully change from time to time, and provides insights for researchers that are interested in collecting multi-session stop-signal task data.

Stopping speed is a state, not a trait

Across the six animals and within each animal, there was a lot of variability in stopping times. Before we quantified how variable and dissimilar they were to each other, we examined whether consecutive single-session SSRTs were following a consistent logical trend. Trend analysis revealed that SSRTs were not consistently following a pattern, across animals and within-animal across stop-signal sides. Because changing SSRTs were not simply explained by something like an increasing or decreasing trend, we determined how

dissimilar they were as compared to the across-session SSRT. We found that more than half of the single-session estimates' confidence intervals were not coinciding with the across-session SSRT, only a minority of single-session SSRTs were likely to share the same SSRT with other sessions, and neighboring sessions had only slightly more similar stopping times. Trends in SSRTs only marginally explained these measures, as on average it did not even increase the overlap between within-session SSRTs and across-session SSRTs with 10%. These findings suggest that most single-session SSRT estimates are not only numerically, but also statistically different from an across-session, trait-like SSRT. Altogether, this supports the hypothesis that stopping speed is a state-like characteristic that meaningfully changes over time. Previous literature has already shown that within-subject stopping speeds can be different when task designs are different (Gordi et al., 2019; Doekemeijer et al., 2023; Weber et al., 2024), but to our knowledge this is the first report that rigorously demonstrates to what degree and partly why stopping speeds change under identical experimental conditions.

SSRT reliability correlates with different task parameters

Sessions with higher variability in go trial response times were associated with lower levels of SSRT reliability. This could be explained by the possibility that the tracking procedure of stop-signal delays is less effective in such sessions. If response speeds are varying on the order of 100s of milliseconds across trials, but the stop-signal delay is limited to adaptations of 50 ms after each stop trial, it may be that this negatively impacts the reliability of the SSRT estimate. For example, an SSD of 350 ms in one trial may be very easy in a trial with slow going speed, but disproportionately more difficult if the SSD changes to 400 ms while the rat goes 300 ms faster the next stop trial. This could make the tracking procedure less effective, in turn affecting the reliability of the SSD used for estimating the SSRT.

Increased skewness of the go trial response time distribution is thought to affect SSRT reliability negatively (Verbruggen et al., 2019), as it may reflect slowing behavior over the course of a session in an attempt to outsmart the tracking procedure (which increases the skewness of the go trial response time distribution), even though this does not pay off with an adaptive SSD (Verbruggen et al., 2013). To our surprise, we found that sessions with higher levels of skewness had more reliable SSRTs. Our speculation is that the estimate of skewness may have been negatively impacted by low numbers of trials contained in the response time distribution, possibly by allowing relatively high response times to have disproportionate effects on skewness,

as this estimate is susceptible to outliers. However, future studies may shed light on whether increased skewness always has positive effects on SSRT reliability in multi-session stop-signal tasks, or whether this relationship is mediated by another unknown factor.

As expected, lower accuracies on stop trials negatively impacted SSRT reliability. The stop-signal task is designed to get as close to 50% accuracy on stop trials as possible to increase the probability of obtaining reliable SSRTs, even when using the integration method (Verbruggen et al., 2019). Although not explicitly tested here, it seems that stop accuracies at the higher end of the spectrum were not necessarily improving the SSRT reliability substantially, which may reflect a ceiling effect. Altogether, across-animal analysis showed that higher go trial response time variabilities, lower levels of skewness in the go trial response time distribution and lower stop trial accuracies negatively impact the reliability of SSRT estimates in our multi-session stop-signal task.

Stop-signal timing and going speed equally contribute to SSRTs

Before we addressed cognitive and neural mechanisms that could possibly drive changing SSRTs, we investigated whether stop-signal timing and going speed were equally contributing to SSRT estimations or not, as they are both variables that make up the SSRT by subtracting the session SSD from the n -th go trial response time that matches 50% performance in that session. Although one might intuitively think that they are contributing equally by definition due to simple subtraction, this is not necessarily the case because one variable may vary substantially more from session to session than the other. For some animals variance in one variable explained more SSRT variance than the other variable, but for other animals this was the opposite. However, across animals, we statistically and visually demonstrated that variance in SSD and n thRespTgo equally contributed to varying SSRTs. Sessions in which animals were going slower were not always accompanied with equally delayed stop instructions, meaning they were slower stoppers in those sessions. Similarly, sessions in which animals were going faster were not always accompanied with equally earlier stop instructions, meaning they were faster stoppers in those sessions. This again supports the notion that SSRTs are not constant, but change from session to session.

Stopping speed may vary with changing motivation, going speed and attention

When animals were highly motivated to get a reward quickly (as indicated with low release times on go trials), for example because they were thirsty, they

also had faster stopping speeds. However, this relationship was observed for sessions with early stop-signal presentations (low SSDs), and to a lesser degree in sessions where stop-signals were presented relatively late. This implies that rats were mostly stopping faster during highly motivational states in sessions where stop-signals were presented fairly early, but not when presented late. When comparing this to the traffic light example, this would mean that you are well capable of stopping quickly when you are in a rush to reach your destination if the traffic light turns red quite distant from the intersection. But, in comparable rushing circumstances, you turn out to be a slower stopper when it turns red when you are closer to the intersection. Why would this happen? And specifically, why would proximity to the intersection impact stopping speed? We speculate that highly motivated individuals possess some degree of impatience that affects their sensitivity to late-presented stop-signals. Some time after starting to go they may just decide (consciously or unconsciously) to not process new incoming visual stimuli, as they focus solely on getting a reward as quickly as possible. This cognitive tunneling may specifically slow down stopping speed in sessions with late-presented stop-signals, due to increasing cognitive load the later a stop-signal is added to the already focused state of mind. Another thought could be that sessions in which rats release early from the central port have more likelihood of stop-signals being presented late in the go-process, increasing the likelihood of failing at stopping in time. However, in this situation the SSD would decrease in turn to compensate for poor performance on stop trials. Moreover, in sessions where the average SSD and RTgo are relatively low, the positive relation between motivation and stopping speed is strongest.

For some animals, faster going speeds were associated with faster stopping speeds. We want to stress that this effect was not present for every animal, implying that some animals support the hypothesis of shared motor dynamics between going and stopping, while other animals do not support this hypothesis. However, no animal's data supported the hypothesis that going and stopping speeds were negatively correlated. In chapter 3 of this dissertation (Figure 2C), we also investigated this relationship and included more animals and more sessions. Here we did observe a significant positive correlation between going speed and stopping speed across animals and trial sides (left/right). Future multi-session stop-signal task studies have to elucidate whether this relationship is replicable.

Sessions in which animals were highly attentive, as indicated by fast release times and high response accuracies on go trials, were accompanied with

faster stopping times as compared to sessions where animals displayed signs of lower attentive states, as approximated with slower go trial release times and response accuracies. As with many cognitive tasks, adequate attention to a relevant stimulus is key for trial outcome success. Although *attention* could have many definitions, it is generally associated with better performance in combination with faster reaction times (Carlson et al., 1983; Prinzmetal et al., 2005). In our task, rats had to divide their attention to both the left and right visual field to be able to quickly release from the central port after go-signal onset and respond correctly. Response accuracy would only reach high levels when release from the central port (marking RTgo) was not just triggered by non-spatially noticing the go-signal, but by being attentive to the spatial location of the go-signal. Our data suggest that rats are also more attentive to a potential stop-signal when they are (spatially) attentive to the go-signal. Translated to a real-world human example this would mean that red traffic lights can be presented later if you are anticipating the emergence of the stop-signal, and still be able to stop in time because your stopping speed increases with increasing attention to the red traffic light.

Other possible explanations for changing stopping speeds

One could argue that the level of training or developmental age may play a role in changing stopping speeds. However, trend analysis of SSRTs demonstrated that stopping times did not consistently decrease with time across animals (Figure 6). Also, our animals reached a certain training level before they underwent surgery for electrode implantation. After that, the behavioral data were collected, so there is no data available with which we can check whether fairly young rats with not fully developed frontal cortices are slower at stopping than fully-developed rats.

Another rationale would be that a speed-accuracy trade-off affected stopping speeds. However, when rats decide to slow down to increase the likelihood of responding correctly on stop trials (although it results in longer waiting times for a reward), our adaptive procedure automatically tunes the stop-signal delay to a level where they will have equally difficult times with stopping as in sessions where they are moving faster. Successful stopping is followed by a later onset of the stop-signal, meaning it becomes more difficult to stop in time, while failed stopping is followed by an earlier onset of the stop-signal, resulting in a higher probability of stopping in time. On the contrary, when rats decide to increase their speed to shorten the waiting time of receiving a reward (while diminishing the probability of getting one), this will also tune

the stop-signal delay to equally difficult levels as a result of how well the rat performs on stop trials. For go trials, adjusting the going speed is rewarding in our stop-signal task, as slowing down simply decreases the likelihood of making erroneous responses (unless they are inattentive). But for stop trials, adjusting going speed is not rewarding due to the adaptive nature of the stop-signal delay based on stopping performance. In fact, we cannot think of a reason why favoring speed or accuracy would impact stopping speed in particular. One could argue that going faster makes stopping more difficult, as there is a cognitive emphasis on going rather than stopping, but we already learned that going faster is not associated with stopping slower.

Limitations and future directions

The task we used was not optimized for extracting which factors contribute to changing SSRTs and SSRT reliability, as it was designed for obtaining as reliable SSRT estimates as possible. Nevertheless, with our analyses that were mostly correlative of nature, we attempted to shed some light on the statistical significance and reliability of single-session SSRTs. This study could be a starting point for researchers that are interested in using multi-session stop-signal task designs to gain more insights about reactive stopping.

However, the sample size of six animals, while providing initial insights, limited statistical power and the ability to generalize findings. Future studies should consider bigger sample sizes to validate these preliminary findings. Moreover, the relationships between factors that are associated with SSRT reliability remain complex and somewhat contradictory to previous literature. While we speculated that low trial numbers might have disproportionately affected skewness, studies with larger trial numbers and experimentally controlled conditions are needed to clarify these relationships. Another limitation lies in the exploration of cognitive mechanisms underlying SSRT variability. While we identified potential contributions from motivation, shared motor dynamics, and attention, these factors were not directly manipulated or measured in a way that allowed us to make causal inferences. Future studies should incorporate more direct assessments of these cognitive states, perhaps through manipulations, to better understand their influence on stopping speed.

In conclusion, while our findings demonstrate that stopping speeds are not fixed within-animal, future research with larger sample sizes, more controlled experimental conditions, and direct measures of cognitive states will be essential for advancing our understanding of reactive stopping in the context

of multi-session approaches. Despite these limitations, we can solidly argue that in a multi-session stop-signal task design, researchers should use *single-session* SSRT estimates instead of an *across-session* SSRT estimate, as within-animal stopping speed substantially and meaningfully changes across sessions. This could turn out to be very useful when the stop-signal task is combined with neural recordings such as local field potentials, because single-session SSRT estimates allow one to time-lock to the SSRT of that session specifically to extract stop-related neural signatures in a more time-precise manner.

Chapter 3

Decreased beta power and OFC-STN phase synchronization during reactive stopping in freely behaving rats

Published as:

Ter Horst, J., Boillot, M., Cohen, M. X., & Englitz, B. (2024). Decreased beta power and OFC-STN phase synchronization during reactive stopping in freely behaving rats. *The Journal of Neuroscience*, 44, 1-13, doi:10.1523/JNEUROSCI.0463-24.2024

Abstract

During natural behavior, an action often needs to be suddenly stopped in response to unexpected sensory input - referred to as reactive stopping. Reactive stopping has been mostly investigated in humans, which led to hypotheses about the involvement of different brain structures, in particular the hyperdirect pathway. Here, we directly investigate the contribution and interaction of two key regions of the hyperdirect pathway, the orbitofrontal cortex (OFC) and subthalamic nucleus (STN), using dual-area, multi-electrode recordings in male rats performing a stop-signal task. In this task rats have to initiate movement to a go-signal, and occasionally stop their movement to the go-signal side after a stop-signal, presented at various stop-signal delays. Both the OFC and STN show near-simultaneous field potential reductions in the beta frequency range (12-30 Hz) compared to the period preceding the go-signal and the movement period. These transient reductions (~200 ms) only happen during reactive stopping, which is when the stop-signal was received after action initiation, and are well-timed after stop-signal onset and before the estimated time of stopping. Phase synchronization analysis also showed a transient attenuation of synchronization between the OFC and STN in the beta range during reactive stopping. The present results provide the first direct quantification of local neural oscillatory activity in the OFC and STN and interareal synchronization specifically timed during reactive stopping.



Introduction

Being able to stop previously initiated actions is crucial for successful behavior in a constantly changing environment. Consider for example a traffic light suddenly turning red while you are about to cross the road. Humans and animals are naturally able to cease their motor plan following external input, suggesting that the underlying brain circuitry is able to quickly terminate a prepared and already initiated movement. This form of inhibition, here referred to as *reactive stopping* or simply *stopping*, is driven by an external stimulus and usually needs to be initiated instantaneously due to the urgent nature of stopping actions, as a failure to stop may cause collision with other objects or leads to other unwanted outcomes.

Reactive stopping in humans is often studied using the stop-signal task, where participants initiate a movement in response to a go-signal, and an occasional stop-signal appears instructing them to stop the initiated action (Logan & Cowan, 1984; Verbruggen & Logan, 2009; Verbruggen et al., 2019). The timing between the go-signal and the stop-signal – the stop-signal delay (SSD) – is critical for successful stopping: as the SSD increases, the probability of stopping before action completion decreases. This is analogous to the increased difficulty of stopping at a red traffic light the closer you are to the intersection. Next to human studies, there is an increasing number of studies using rodents to investigate reactive stopping, especially because they enable researchers to study anatomy, neurotransmitters and neural activity in more detail.

There is converging evidence showing that the human right inferior frontal cortex, the pre-supplementary motor area and the subthalamic nucleus (STN) are key brain areas involved in stopping (Aron et al., 2004; Aron et al., 2007; Aron et al., 2014; Aron et al., 2016; Jahanshahi et al., 2015). A recent study in humans showed that beta-band power (12–30 Hz) in the right inferior frontal cortex precedes the pre-supplementary motor area, and predicts stopping performance (Schaum et al., 2021). Similar to those frontal brain areas, studies reported increases in STN activity during stopping (Kühn et al., 2004; Swann et al., 2011; Zavala et al., 2015). However, studies that investigated right inferior frontal cortex and STN activity are not consistent, as some reported beta modulations after stopping has occurred (Hubbard & Sahakyan, 2023; Swann et al., 2009; Swann et al., 2011), and some before (Swann et al., 2012; Wagner et al., 2018; Wessel et al., 2016). The so-called hyperdirect pathway is

hypothesized to connect the right inferior frontal cortex with the STN (Aron et al., 2007; Chen et al., 2020; Haynes & Haber, 2013; Nambu et al., 2002; Magill et al., 2004), suggesting that stopping may be facilitated by this pathway.

Taken together, these studies suggest that functional connectivity between the inferior frontal cortex and STN may play a role in reactive stopping, but evidence for this during actual execution of stopping is sparse (Aron et al., 2016) and the timing of beta power modulations are inconsistent across studies. Different species, recording methods, behavioral tasks and criteria for which activities are stop-related could partially explain these inconsistent findings. In addition, recording deep brain structures in humans is challenging, and recording from patients makes it difficult to dissociate normal activity from disease-related activity. Recordings from inferior frontal cortex either have poor temporal precision (functional magnetic resonance imaging) or lack anatomical precision (electroencephalography). These methodological difficulties and limitations can be addressed using rodent models.

Here, we aimed to record brain activity in the rat orbitofrontal cortex (OFC) and STN at high temporal and anatomical precision during the stop-signal task. The rat OFC and STN are functionally comparable to the human inferior frontal cortex and STN (Eagle et al., 2008; Eagle & Baunez, 2010; Parent & Hazrati, 1995). However, caution is warranted because functional homologs may not always match anatomical homologs (Robbins, 1998). We show that the OFC and STN decrease their beta power and interareal phase synchronization during stopping, specifically after the stop-signal but before the estimated stopping time.

Materials and Methods

Animals

Ten male wild-type Long-Evans rats participated in this study, aged 9 weeks and weighing 250-320 grams at the start of behavioral training (Charles River Laboratories, Calco, Italy). Rats were housed pairwise in Makrolon type III cages (UNO B.V., Zevenaar, The Netherlands) with a reversed 12-hour day-night cycle in a temperature- and humidity-controlled room ($21 \pm 2^\circ\text{C}$, $60 \pm 15\%$). As soon as the rats weighed more than 350 grams, they were housed in Makrolon type IVS cages to provide more horizontal space. As soon as the rats were implanted with electrodes they were housed individually, and the low conventional cage lid was replaced by a high cage lid to prevent damage

to the implant. Corn cob granules were used as cage bedding, and sizzle bedding and a cardboard shelter were provided as cage enrichment. The rats were put on a restricted water intake schedule as soon as they acclimatized in the research facility. Every Monday to Friday morning the rats could get water in the behavioral task (~5-8 mL, depending on performance), and in the afternoon they could drink *ad libitum* water for 30 minutes from a bottle. During weekend days, the rats received 30 grams of hydrogel (ClearH2O Inc., Westbrook, Maine, USA) each day, to keep the daily intake of water as stable as possible. Food pellets were provided *ad libitum* at all times. Weight and health were monitored on a daily basis. All animal procedures were approved by the Animal Welfare Body of the Radboud University Nijmegen and the Animal Experiment Committee (CCD No. AVD10300 2016 482, Project No. 2015-0129), according to national and international laws, to protect welfare under experimental conditions.

Skinner box

After acclimation of two weeks in the research facility the rats started with the restricted water intake schedule and behavioral training. Training and testing took place in a custom-built Skinner box (inside dimensions: 25 × 27 × 25 cm), with one wall containing three nose-poke ports (bottom-left, bottom-center, bottom-right, see Figure 1A/B). Each port had an infrared emitter and phototransistor (type L-53F3C and L-53P3C, peak 940 nm, Farnell B.V., Utrecht, The Netherlands) enabling continuous automatic detection of a nose-poke by the rat. In addition, the left and right port also contained green light-emitting diodes (type L-53SGD-5V, peak 565 nm, Farnell B.V., Utrecht, The Netherlands) for presenting visual stimuli, as well as a small silicone tube at each bottom of the port for providing 50 µL water drop rewards driven by solenoid pumps (The Lee Company, Westbrook, Connecticut, USA). The roof of the Skinner box contained an entry for the recording cable and a simple USB camera was mounted for collecting video frames during the behavioral task (25 Hz frame rate, 720x575 resolution). Electronics needed for the task in the Skinner box were controlled by a computer with custom-written code in MATLAB (R2018b, The MathWorks Inc., Natick, MA, USA).

Training procedure stop-signal task

The rodent version of the stop-signal task was based on the behavioral tasks used by Feola et al. (2000) and Bryden et al. (2012), with optimizations taken from Verbruggen et al. (2019). The training procedure consisted of three main steps, namely 1) central nose-poke initiation, 2) unilateral cue discrimination,

and 3) stop trial introduction. Rats were trained for maximally one hour or 200 trials each day, while having rest during weekend days. During the central nose-poke initiation phase, the rats had to learn to initiate a trial by poking their nose in the central port. On the first day of training, poking at the central port for the shortest time detectable was already enough to initiate a trial, causing the presentation of a light cue (go-signal) for 100 ms either on the left or right side with an immediate water drop reward provided at the corresponding side. As the reward was provided immediately after trial initiation the rat was not required to make a correct response yet. Despite not needing a correct response yet, trials were never terminated before a response was recorded, causing the rats to learn that they always had to respond in either the left or right port before they could initiate a new trial. This response requirement was kept at all training phases and data acquisition. The idea of this phase was to let the rat learn to associate central nose-poking with a positive outcome. Each day, the time required to stay in the central port (i.e., trial initiation time) was increased by 100 ms if the rat managed to initiate at least 100 trials in the previous session. When rats would release their nose from the central port before the required trial initiation time, the trial was not initiated and had to be re-initiated completely. This means premature responses as described in Verbruggen et al. (2019) are not possible. As soon as rats could initiate at least 100 trials in a session with 1000 ms initiation time, they proceeded to the second phase of training.

In the second phase, the rats had to learn to respond correctly to the go-signal before getting a reward. In practice this means the rat initiated a trial, received a go-signal at either the left or right side, and would only get a reward when the rat poked his nose in the port where the light was presented. This allowed the rats to learn to associate the go-signal side with the reward side. When the response was correct, the reward was immediately provided at the response port. No reward was provided after responding incorrectly. During this phase of training, the minimal time between the response at the lateral port and initiating a new trial (inter-trial interval) was gradually increased from 0 to 3 s in steps of 500 ms each session until 3 s was reached, to allow for proper separation of trials and to prevent rushing. To prevent go-signal anticipation, a jitter was slowly added to the trial initiation time from 0 to ± 200 ms in steps of 50 ms every following session until a jitter of ± 200 ms was reached. For each trial, the jitter value was randomly selected from a uniform distribution of numbers ranging from -200 to 200 ms with steps of 10 ms. As soon as the rats reached response accuracy above 80% for at least five consecutive days, they moved to the third and final phase of training.




Figure 1. Dual-area, multi-electrode recordings in the OFC and STN during reactive stopping. **A**, Rats were trained on a stop-signal task where they initiated trials autonomously, with 75% go trials and 25% stop trials. Correctly performed trials were rewarded with a water drop. Stop trials were segregated based on stop-signal timing relative to release and correctness of response, as indicated by the different colored boxes. **B**, Trained rats were implanted in the right OFC and STN (32 electrodes in the OFC, 30 electrodes in the STN) and could freely move in the Skinner box with the use of a commutator and lever arm. **C**, Custom-designed probes whose geometry was adapted to the anatomical shapes of the OFC and STN (left). Rectangles and parallelograms in coronal sections represent electrode positions for each animal recovered from histology. Dorsal view and spacing of electrodes (right top). Representative raw data shown from three electrodes (right middle). Baseline power shown for each animal in different colors (right bottom). Coronal sections taken from Paxinos and Watson (7th ed.). RT = release time; MT = movement time; RespT = response time; SSD = stop-signal delay; DV = dorsoventral axis.

In the final phase, stop trials were slowly introduced in addition to the go-only trials. On the first and second session of this final training phase the rats received 10% stop trials, the next three sessions 20% stop trials, and from the sixth day onwards 25% stop trials. In stop trials, the go-signal was followed by a stop-signal after a variable stop-signal delay (SSD). The stop-signal was given by a light on the other side than the go-signal side, and stayed illuminated until the rat made a response at either the left or right port (Figure 1A). The SSD was determined separately for left and right stop trials through a staircase procedure, where it increased with 50 ms in case of a correct response, and decreased with 50 ms in case of an erroneous response. The SSD was not separately determined for two animals, and for two other animals only after ~75% of their sessions were collected. The starting SSD for stop trials was determined by averaging the SSDs from all stop trials from the previous session, for left and right stop trials separately. This staircase procedure is standard in human studies, ensures the collection of a wide range of SSDs and helps with obtaining a reliable stop-signal reaction time (SSRT) estimate (Verbruggen et al., 2019). Having a separately determined SSD for left and right stop trials ensures that the difficulty of a stop trial is always comparable between left and right stop trials, as a possible response bias may lead to an imbalanced accuracy for left versus right stop trials when there is a shared SSD. Go-signal side and stop trial occurrence were randomized and balanced as such that in each set of eight trials three left go trials, three right go trials, one left stop trial, and one right stop trial were randomly shuffled. Response bias was continuously checked by computing the percentage of left and right responses in the last 20 trials. When one side fell below 35%, the next trial was replaced by a go trial to that side to discourage response bias. As soon as the accuracy on stop trials floated around 50% and go trial accuracy was above 80% for five consecutive days, the rats were ready for electrode implantation. Training took approximately 6-8 weeks (30-40 sessions), and all rats achieved these criteria and were included in the subsequent recordings.

Custom-designed probes

For recording intracranial local field potentials in the orbitofrontal cortex (OFC) and subthalamic nucleus (STN), custom-designed probes were built (see Figure 1C for an illustration and França et al., 2020). As the anatomical structures of interest have their own unique shape and size, the length of the individual electrodes was designed based on a rat brain atlas (Paxinos & Watson, 2013). Each probe consisted of 50 μm diameter tungsten wires (99.95%, California Fine Wire Company, Grover Beach, CA, USA), covered in

a styrene-isoprene-styrene polymer except for the tip, allowing for recording local field potentials at the tip of each wire. The OFC probe was built with an anterior to posterior downward angle, whereas the STN probe was built with a lateral to medial downward angle. With the help of three custom-designed perforated grids (Eurocircuits, Mechelen, Belgium) placed in series, wires were aligned in parallel to each other (see Figure 1C for a dorsal view of the design), whereafter the depth of each row of wires was adjusted with the use of a caliper. Core-to-core distance between wires was $\sim 250\text{ }\mu\text{m}$ in the horizontal plane. The wire configuration was glued together with OptiBond (Kerr, Kloten, Switzerland) and in turn glued onto a custom-designed printed circuit board (PCB; Eurocircuits, Mechelen, Belgium). An Omnetics connector (type A79026-001, Omnetics Connector Corporation, Minneapolis, MN, USA) was soldered onto the PCB, allowing for a 32 channel Intan headstage connection. Each tungsten wire was soldered to a single contact that led to a single channel on the headstage. In addition, a $51\text{ }\mu\text{m}$ diameter stainless steel ground wire coated with perfluoroalkoxy alkane insulation was soldered to the PCB (Science Products, Hofheim, Germany), which was connected to a screw over the cerebellum during surgery. A stainless steel screw with 3.2 mm head diameter (M1.6 x 4 mm, RS Components B.V., Haarlem, The Netherlands) was glued to each upper corner of the PCB with epoxy, as such that a rubber band could easily be placed behind the screw head to secure the cable and headstage in the implant while the rat was in motion during the task. The entire PCB was covered with epoxy for protection and electrical insulation. Channel impedance was on average $181\text{ k}\Omega$ (± 66). For more probe manufacturing details, see França et al. (2020).

Implantation surgery

As the rats were on a restricted water intake schedule during training, they were provided with daily hydrogel five days before surgery to ensure hydration before surgery. They received 25 mL of hydrogel five days before surgery, about two hours after training, and this amount was gradually increased on the following days until 60 mL hydrogel the day before surgery. On the day of surgery, rats were anesthetized with 5% isoflurane in a mixture of oxygen (0.25 L/min) and air (0.5 L/min). Rats were shaved on the dorsal surface of their head and received a subcutaneous injection of carprofen for analgesia (Rimadyl, 5 mg/mL, 1 mL/kg body weight). Once placed in a stereotactic frame (Kopf Instruments, Los Angeles, CA, USA), isoflurane was set back to 2% for the remaining part of surgery to maintain appropriate depth of anesthesia. Heart rate and oxygen saturation were continuously monitored during surgery,

and temperature was maintained around 37 °C with a closed-loop heating plate and rectal temperature probe. Every passing hour, a 2 mL subcutaneous injection of saline was given to support hydration during anesthesia. The incision site was disinfected with betadine and subcutaneously injected with 0.25 mL of a 2:1 lidocaine-bupivacaine mixture for local anesthesia (Lidocaine HCl, 10 mg/mL, 25 µL; Bupivacaine Actavis, 5mg/mL, 125 µL; diluted with 0.2125 mL NaCl). A midline incision was made after which 30% H₂O₂ was used to clean the skull and easily visualize bregma and lambda. Head tilt was corrected until the tilt difference was at least lower than 200 µm. Four small countersunk screws (DIN 965 A2 304 M1x2, Screws and more GmbH, Ennepetal, Germany) were placed in the skull for dental cement to bond to at a later stage of surgery, and craniotomies were made for the right OFC and right STN, based on their atlas locations relative to bregma (Paxinos & Watson, 2013). The center coordinate of the OFC wire tips was AP 3.2 mm and ML 2.3 mm relative to Bregma, and DV -5.0 mm from dura surface, whereas the STN wire tips had a center coordinate of AP -3.6 mm and ML 2.6 mm relative to Bregma, and DV -7.7 mm relative to dura surface (Figure 1C, gray-colored crosses in dorsal array view). The probes were lowered very slowly and fixed with dental cement once they reached final depth (Super-Bond C&B, Sun Medical, Moriyama, Japan). Next, the ground wire of each probe was wrapped around the thread of a countersunk screw (DIN 965 A2 304 M1.6x3, RS Pro GmbH, Frankfurt am Main, Germany) over the cerebellum which reached dura matter. The wired screw was covered with silver conductive lacquer (RS Components B.V., Haarlem, The Netherlands) to establish a proper connection. After the silver lacquer was dry, a final layer of dental cement was applied to cover all screws and strengthen the entire implant (Paladur, Kulzer GmbH, Hanau, Germany). When this was cured, isoflurane was slowly decreased to 0%, and 100% oxygen (0.75 L/min) was given for five minutes to boost blood oxygenation while the rat was exiting anesthesia. A heating pad was provided under half of the cage for the upcoming night, as well as moistened booster food (Ssniff Spezialdiäten GmbH, Soest, Germany), regular food pellets and 60 grams of hydrogel (ClearH₂O Inc., Westbrook, ME, USA). About 22 hours post-surgery, rats received another subcutaneous injection of carprofen for analgesia (Rimadyl, 5 mg/mL, 1 mL/kg body weight). After three days of recovery, moistened booster food was not provided anymore and the amount of hydrogel provided was slowly decreased and replaced by the restricted water intake schedule. One week post-surgery rats were performing the task again on a daily basis, and two weeks post-surgery local field potential recordings were started.

Data acquisition

During the workweek rats performed the stop-signal task daily for one hour or maximally 200 trials, while intracranial local field potentials were recorded from the OFC and STN at a sampling rate of 30 kHz. In addition, video frames were collected with a frame rate of 25 Hz and aligned with the electrophysiological data by using the task event markers. For both the OFC and STN a 32-channel recording headstage was used to amplify the signals, whose outputs were merged using a dual headstage adapter, that in turn connected to a single RHD SPI interface cable (Intan Technologies, Los Angeles, CA, USA). This interface cable was cut at approximately 40 cm above the dual headstage adapter, where a miniature commutator (Adafruit, New York, NY, USA) was soldered in between both ends of the interface cable, which allowed the rats to rotate in the horizontal plane as much as they wanted. The cable arrangement went through the roof of the Skinner box, where a simple lever system with counterweight was used to allow for low-resistance movement in the vertical plane (Figure 1B). Rats could easily rotate their body and move up and down with their head without needing effort to displace the cable while executing the stop-signal task. The interface cable was connected to an acquisition board (Open Ephys, Atlanta, GA, USA). This acquisition board also received task-event markers from the computer controlling the Skinner box, which were displayed together with the electrophysiological data and saved to disk on a separate acquisition computer.

Perfusion and histology

After data acquisition, rats were euthanized and transcardially perfused. For fixation purposes, a 4% solution of paraformaldehyde (PFA) in phosphate-buffered saline 1x (PBS) was prepared the day before. When needed, pH was adjusted to 7.2 with either sodium hydroxide or hydrochloric acid. PFA 4% was kept at 4 °C, as well as a PBS 1x solution. On the day of transcardial perfusion, rats were anesthetized with 5% isoflurane in a mixture of oxygen (0.25 L/min) and air (0.5 L/min) and received an intraperitoneal injection of pentobarbital (60 mg/mL, 6 mL/kg body weight). After a final check for absence of reflexes the thoracic cavity was exposed. First, PBS kept at 4 °C was transcardially pumped through with a flow rate of 20 mL/min for 10 minutes. Then, liquid flow was switched to PFA 4% kept at 4 °C with a flow rate of 20 mL/min for 10 minutes. When the perfusion procedure was finished, the head was kept in PFA 4% at 4 °C for 24 hours. Next, the brain was extracted and again kept in PFA 4% at 4 °C for 24 hours. Until sectioning, the brain was stored in PBS 1x with 0.01% sodium azide at 4 °C. A Leica VT1000 S vibratome was used for

making 60 μm coronal sections. Sections were stained with 4',6-diamidino-2-phenylindole (DAPI, Thermo Fisher Scientific, Waltham, MA, USA) for 10 minutes in well plates, and two times transferred to clean well plates with fresh PBS 1x and washed for 10 minutes. Coronal sections were mounted on gelatine-coated slides, air-dried, and covered with FluorSave (Merck Life Science N.V., Amsterdam, The Netherlands) and coverslips. Slides were kept in the dark at -20°C until fluorescent images were made with a Zeiss Imager A2 microscope. Electrode tracks in tissue were imaged with a 2.5x objective, CCD camera and ZEN 3.2 software to check for implant accuracy. For all rats, the electrodes were in or in close vicinity of the target regions (Figure 1C, colored zones), except for the STN electrode array from one animal. The results obtained from this animal were nonetheless qualitatively well matched and were therefore still included in the analysis.

Behavioral data analysis

Analyses of behavioral data were done with custom-written code in MATLAB software (R2018b, The MathWorks Inc., Natick, MA, USA). For each trial, release time was defined as the time between go-signal onset and release from the central port, movement time was defined as the time between release from the central port and response at either the left or right port, and summed together they form response time. Importantly, response time is the homologue of reaction time in human stop-signal tasks, as in human tasks reaction time is defined as the time between go-signal and response. In addition, for stop trials the stop-signal delay (SSD) was defined as the delay between go-signal onset and stop-signal onset (Figure 1A). To deal with anticipatory releases (although discouraged with a jitter for go-signal onset), a local minimum in the smoothed bimodal release time distribution was identified and used as the lower bound for trial removal, as this portion of releases was anticipatory instead of reactive and attentive to the go-signal. The maximal lower bound was set at 100 ms, and the upper bound was set at three standard deviations above the mean release time. As no time limit was set for movement times, trials with movement times above five seconds were removed initially, whereafter the upper bound was set at three standard deviations above the mean for the remaining trials. The remaining trials were used for computing average release times, movement times, response times, stop-signal delays and stop-signal reaction times per session. The unobservable stop-signal reaction time was estimated with the integration method (Verbruggen et al., 2019) for each session separately. In short, response times on go trials were sorted from fastest to slowest, and the n -th response time corresponding

to the point where the integral is equal to $p(\text{respond}|\text{signal})$ was found by multiplying $p(\text{respond}|\text{signal})$ from that session with the number of go trials in that session. Next, the SSRT was computed by subtracting the average SSD of that session from the n -th response time. As the SSD was separately determined for left and right stop trials, the SSRT was also separately estimated by using left and right response time distributions and SSDs. The response time distribution did not contain replacements for go omissions (go trials without a response), as responses were required on all trials. In addition, premature go responses (in our task releases before go-signal onset) are not part of the response time distribution either, as trials are not initialized in those cases (see training procedure). Both correctly and incorrectly performed go trials were used for the distribution. Accuracy on go and stop trials was computed by dividing the sum of correct responses on go and stop trials by the total number of go and stop trials, respectively. To investigate how release time, movement time, and response time were affected by trial type (go/stop) and response outcome (correct/incorrect), a 1-way ANOVA was used, and the interaction between trial type and response outcome was analyzed with a repeated measures ANOVA. Relationships between response times on go trials, SSD and SSRT were estimated with Spearman's correlation coefficient on a session-by-session basis.

Stop trial segregation for extracting the true stop-related activity

As the SSD could also be timed as such that the stop-signal was presented before release at the central port, we segregated stop trials with a correct response outcome from stop trials with an incorrect response outcome based on whether the stop-signal was presented after or before release. This allowed for dissociating power dynamics after instructed stopping while movement to the go-signal side was ongoing from when the animal was not yet moving to the go-signal side. This segregation resulted in four different stop trial conditions with their own color coding in Figure 1, 3 and 4: (a) correct responding in trials when the rat was still in the central port during stop-signal onset, practically meaning the rat was able to go to the stop-signal side without actually needing to physically stop movement to the go-signal side - blue; (b) incorrect responding in trials when the rat was still in the central port during stop-signal onset, representing those trials where the rat erroneously went to the go-signal side despite not having to stop physical movement to the go-signal side yet - yellow; (c) incorrect responding in trials when the rat was moving during stop-signal onset, where the rat already left the central port during stop-signal onset towards the go-signal, but failed to stop the

ongoing movement towards the go-signal side before making the erroneous response (unsuccessful stopping) - red; (d) correct responding in trials when the rat was moving during stop-signal onset, and the rat correctly stopped the ongoing movement to the go-signal side and moved to the stop-signal side (successful stopping) - green. The latter is the only stop trial condition where physical reactive stopping truly occurs, while in the first stopping condition the rat at most had to stop the plan to go to the go-signal side. As go trials are traditionally segregated into fast and slow go trials for comparison with unsuccessful and successful stop trials, respectively, we also segregated them this way. Go trials were segregated into slow or fast go trials based on the response time relative to the average SSD and estimated SSRT; go trials were considered slow when the response time was greater than the sum of SSD and SSRT, and considered fast when it was smaller. This was done for each session separately. Across animals, 52.1% of go trials were marked as slow (95% CI [49.8 54.4]), and 47.9% as fast (95% CI [45.6 50.2]). Among stop trials, 26.0% of trials were marked with a correct response while the stop-signal was presented before moving (95% CI [20.4 31.7]), and 23.5% during moving (95% CI [17.6 29.5]), while 8.7% of trials were marked with an incorrect response while the stop-signal was presented before moving (95% CI [4.0 13.3]), and 41.8% during moving (95% CI [38.1 45.5]).

Although one would intuitively think that successful stopping should be compared to unsuccessful stopping, we considered this contrast problematic because we observed a beta power reduction related to arrival at the lateral port that would confound the contrast as it would imply we compare two different behavioral events (stopping vs. arriving at port) when keeping the traditional stop-signal alignment. Alternatively, contrasting successful stopping with slow go trials is arguably another good contrast because stop trial activity gets rid of activity that is related to the going process, assuming this activity is emerging in the OFC and STN. We decided to compare successful stopping with the stopping condition where stopping was instructed, but did not result in stopping, as there was no ongoing movement to be stopped. For these two contrasted conditions the stop-signal onset was almost equally distant to the response at the lateral port, resulting in aligned stop-signal-related processing. The only element that was obviously different, was that in the successful stopping condition the go process was physically ongoing during stop-signal onset, while it was not for the other condition. However, this could not confound the results, as there were no indications that the release and ongoing go-process were emerging in beta power changes in the OFC and

STN. In addition, when successful stopping would have been contrasted with the other stopping condition (stop-signal onset while moving, but incorrect response), conclusions would have been the same as this condition showed a comparable beta power envelope.

Local field potential data preprocessing

Preprocessing was done using MATLAB software with a combination of custom-written code and the EEGLAB toolbox (Delorme & Makeig, 2004). Data were downsampled to 1 kHz and high-pass filtered at 0.5 Hz to remove slow drifts. Per session, channels with excessive amplitude artifacts or clearly only containing noise were manually removed. Local referencing was applied by subtracting the common average from each channel time-series for OFC and STN channels separately. Epochs were cut (-2 to +2.5 s relative to go-signal onset) and baseline subtracted (-200 to 0 ms relative to go-signal onset). Trials with poor signal were removed manually, and independent component analysis (ICA) was used to isolate left-over independent components with a noisy source and remove them from the data.

For each animal, the signal was convolved with a set of complex Morlet wavelets (complex sine waves tapered with a Gaussian) to extract time-frequency power for each channel and condition separately, with 20 ms precision and 40 frequencies log-spaced between 1 and 100 Hz. Number of cycles in the complex Morlet wavelets ranged from 3 to 8, log-spaced in 40 steps for the 40 different frequencies. This was equivalent to a time-domain full-width at half-maximum range of 1124 to 30 ms, and a frequency-domain full-width at half-maximum range of 0.5 to 21 Hz. For the beta range of interest (12-30 Hz), the average time-domain full-width at half-maximum was 113.3 ms and the average frequency-domain full-width at half-maximum was 5.86 Hz (Cohen, 2019). Phase synchronization in the beta frequency range was computed for each channel pair and condition separately using inter-site phase clustering (ISPC). Frequencies ranged from 12 to 30 Hz with 20 log-spaced steps, and the complex Morlet wavelets had the same number of cycles as this range of frequencies had for power extraction (i.e., 5.1 to 6.2 cycles, log-spaced in 20 steps). Again, the signal was convolved with a set of complex Morlet wavelets, and the angle differences between channel pairs of this complex analytic signal were computed to find phase synchronization values for each channel pair. Beta power and phase synchronization envelopes were computed by averaging channel activities within-area between 12 and 30 Hz for each condition separately. Average phase synchronization values were computed from -200 to 0 ms relative to SSRT for

plotting channel pair phase synchronization lines in Figure 5C. To facilitate interpretation, we defined a threshold of one standard deviation above the median from all channel pairs and conditions, and interpreted only ISPC values above this threshold. To explore how individual channels of the OFC and STN are connected to the other brain area and act as nodes in this network, we analyzed their “hubness”, a network property known from graph theory that reflects the degree of centrality in a network. Here, we computed hubness for each channel as the proportion of suprathreshold connections it had with all channels from the other brain area over all possible connections with this brain area. This proportion is visually represented using dot size at each channel location. Higher degrees of hubness, reflected with bigger dots, demonstrate that these channels are synchronized with a larger number of channels in the other brain area.

Results

We recorded from freely behaving rats ($N=10$), while they performed a reactive stopping task, in which a go-signal was followed by a stop-signal in 25% of the trials, prompting the rats to stop and reverse their initial movement in order to receive a water reward. Local field potentials from the orbitofrontal cortex (OFC) and the subthalamic nucleus (STN) were recorded with multi-electrode arrays to understand the specific contribution and interaction of these regions in the translation of the stop-signal to a stopping action.

Rats successfully learn to perform reactive stopping

The rats overall performed well on go trials, with an average accuracy of 85.2% (95% CI [83.1 87.3]). Due to the staircase procedure for the SSD, average stop trial accuracy was 50.5% (95% CI [48.7 52.3]) and no rat had a stop trial accuracy below 25% or above 75%. The independence assumption of the horse-race model (Verbruggen et al., 2019) was not violated, as for all rats the average response time on incorrectly performed stop trials was numerically smaller than the average response time on go trials. As the rats did many sessions (on average 36.4), the SSRT was estimated for each session separately to deal with session variations in response speed and SSD (Figure 2B). Across rats, the average SSD was 306.4 ms (95% CI [266.7 346.1]) and the average SSRT was 389.5 ms (95% CI [342.0 437.1]).

Release time, movement time and response time were all not significantly affected by trial type alone (RT: $F(1, 18) = 0.58, p = .456$; MT: $F(1, 18) = 4.07$,

$p = .059$; RespT: $F(1, 18) = 2.41, p = .138$) and response outcome alone (RT: $F(1, 18) = 4.05, p = .059$; MT: $F(1, 18) = 0.10, p = .753$; RespT: $F(1, 18) = 0.27, p = .608$). However, there was a significant trial type $\times$ response outcome interaction on release time, movement time and response time (RT: $F(1, 9) = 119.59, p < .001, \eta_p^2 = .93$; MT: $F(1, 9) = 60.88, p < .001, \eta_p^2 = .87$; RespT: $F(1, 9) = 196.76, p < .001, \eta_p^2 = .96$). Go trials in which the rat released relatively slowly were less likely performed correctly, possibly due to a lack of attention to the go-signal in this subset of trials. Conversely, slow releases in stop trials were associated with increased probability of responding correctly. We interpret the latter by the timing of the stop-signal relative to the ongoing response to the go-signal, as a later release at the central port in stop trials is associated with less proximity to the wrong response port when the stop-signal is presented. This allows for more time to stop as compared to when the rat would be very close to the wrong response port. The interaction effect for movement time is mostly driven by the increased time needed for stopping and moving towards the opposite response port. The interaction effects found on release time and movement time accumulated in the response time interaction: relative larger response times were associated with lower accuracy in go trials, while larger response times were associated with higher accuracy in stop trials (Figure 1A).

Sessions with faster response times on go trials were associated with shorter average SSDs (Figure 2C; $r = .56, p < .001$, Spearman correlation). This can be considered as a speed-accuracy trade-off, as going faster requires convergence to shorter SSDs to still establish 50% accuracy on stop trials. Similarly, sessions with faster SSRTs were accompanied with longer SSDs (Figure 2D; $r = -.16, p = .003$), probably because faster SSRTs allowed for longer SSDs while still maintaining 50% accuracy, although this correlation was relatively weak. Sessions with slower response times on go trials had slower SSRTs (Figure 2E; $r = .63, p < .001$), also reflecting a speed-accuracy trade-off, or suggesting a shared process that controls movement speed and stopping speed, or that both movement speed and stopping speed are jointly limited.

The timing of the SSRT relative to the response matched the response outcome. When stopping succeeded, the SSRT occurred well before the response, as the rat still had to move to the other side to receive a reward. On the other hand, the SSRT occurred after the response when stopping failed, as in these trials the stopping process finished too late relative to the go process for it to be in time for stopping (Figure 3A/B).




Figure 2. Rats perform reactive stopping in the stop-signal task. **A**, Trial type and response outcome alone did not affect RT, MT, and RespT. There was a significant crossover interaction between trial type and response outcome on RT, MT and RespT (see Results for statistics). We used an adaptive SSD to titrate task difficulty per animal, maintaining stop trial accuracy at 50%. Error bars reflect 95% confidence intervals. Asterisks indicate statistical significance ($p < 0.001$). n.s., not significant. $N = 10$ rats. **B**, For each animal, the SSRT was estimated for each session separately, because RespT on go trials and SSD varied across sessions. Each envelope represents the mean across animals per session number for RespT on go trials, SSD and SSRT, respectively. The gray-shaded area represents standard error of the mean. **C**, Faster RespTs on go trials were associated with shorter SSDs. Each dot represents a session, and colors indicate animals. **D**, Faster SSRTs allowed for longer SSDs while keeping 50% stop trial accuracy. **E**, Slower RespTs on go trials were accompanied by slower SSRTs. RespTgo = response time on go trials.

Spectral dynamics of neural activity in the OFC and STN

Average beta power significantly decreased instead of increased after stop-signal onset for both the OFC (Figure 3A) and STN (Figure 3B). This decrease in beta power was apparent in stop trials and in go trials (Figure 3A/B left) and returned back to baseline levels after the response. Simultaneous with the beta power decrease, we observed gamma (35-70 Hz) power decreases in the OFC. This gamma power co-occurred with the observed beta power changes, but seemed to last for a shorter period of time. The STN showed fairly strong theta-alpha (7-12 Hz) power increases that were likely related to the go-process in the task, as the significance appears right after release time (downward arrows) in all conditions and became less strong, or even non-significant, during immobile phases of the task including around the same time as the beta power reductions. In addition, prolonged delta (1-5 Hz) power increases emerged from release at the central port until after the response at the lateral port.

Because stop-related activity should happen after the stop-signal and before the SSRT, and because prior literature indicates an important role for beta-band activity in stopping, we focussed on beta power and phase synchronization in consecutive analyses.


Stopping is associated with beta power suppression before SSRT

The beta power decrease seemed to relate to the response at the lateral port, as for all six conditions the response (Figure 3, upward arrows) happened during or shortly after the peak decrease in beta power. However, we observed an additional decrease in beta power during successful stopping when the stop-signal was presented during ongoing movement to the go-signal side (Figure 3A/B, top right, green outline). This decrease in beta power occurred well before the response at the lateral port, and specifically after the stop-signal but before the SSRT. During stopping, the strongest decrease at 20 Hz contained only 30% of baseline OFC power, while the STN had the strongest decrease at 24 Hz and contained 68% of baseline power. While the initial stop-related decrease in beta power in the OFC was followed by a clear secondary decrease just before the response, STN beta power did not show a clear secondary decrease but a continuation of relatively low beta power and returned to baseline beta power (Figure 4A/B, green envelopes). The observed beta decrease before and around response port arrival found for all conditions could be explained as a stopping signature as well, as stopping is also needed when swiftly moving towards the lateral response port. But, to


properly separate the reactive stopping activity triggered by the stop-signal from the beta decrease related to arrival at the lateral response port, we time-locked the data to response at the lateral port for the beta power envelopes in Figure 4 and the phase synchronization data shown Figure 5.

> **Figure 3.** The OFC and STN show time-locked, transient changes in neural activity during stop-signal task. Each plot shows the local field power in time and frequency, time-locked to the stop-signal for OFC (**A**) and STN (**B**) in the different conditions. As go trials do not have a stop-signal (STOP) and SSRT, their theoretical position given the SSD in that session is indicated by the gray vertical lines (see Materials and Methods). Horizontal black lines in the time-frequency window indicate the beta frequency band boundaries, based on a priori selection (12–30 Hz). Black bars on the time-axis represent the standard error of the mean of go-signal onset (GO), release at central port (downward arrow), and response at lateral port (upward arrow), respectively. Black-colored contours represent cluster-corrected significant power deviations relative to baseline. Conditions are color-matched with beta power envelopes in Figure 4. **A**, The OFC exhibited a response-related beta power decrease just before and around the response at the lateral port, but was also preceded with another decrease in beta power after the stop-signal and before the SSRT, only when stopping was needed and successful (top right, green outline). When the stop-signal was presented while the rats had not released yet, but made the correct response, the early beta power decrease was absent (top middle, blue outline). **B**, The STN also showed a response-related beta power decrease before and around the response at the lateral port, but another beta decrease was present between the stop-signal and SSRT only when stopping was needed and successful (green outline).

A



B



Next, we directly compare the key successful stopping condition (green) with the stopping condition that had most similarity in terms of stop-signal timing relative to response at the lateral port (blue). In the latter condition, stopping was instructed but did not lead to stopping as the stop-signal was presented before the rat released from the central port, leading to a go response to the correct stop-signal side. In both conditions the rat received a stop-signal and responded correctly, but in one condition physical stopping happened, while in the other it did not happen. This comparison demonstrated a significant beta power decrease in the OFC and STN after stop-signal onset but before the SSRT (Figure 4, green-blue markers above the time axis, significance computed using false-discovery rate correction for multiple comparisons). Both the OFC and STN exhibited a beta power decrease almost immediately after the stop-signal with a steep downward slope reaching a group of successive significant timepoints after 180 and 220 ms, respectively, until 13 ms before and 27 ms after SSRT, respectively. During reactive stopping, the deepest level in the beta power envelope contained 39% of baseline power in the OFC and 82% in the STN. Both the OFC and STN showed a significant beta power decrease for about 200 ms, but STN significance was delayed by 40 ms relative to OFC significance. On successful stop trials at a single-trial level, beta power decreased in at least 95% of trials during the pre-SSRT stopping window (-200 to 0 ms relative to SSRT) relative to baseline in the same trial for all animals and both areas (average OFC = 99.4%; average STN = 99.4%; data not shown here).

When stopping was instructed, but there was no movement to be actively stopped as the stop-signal was presented before release from the central port (blue and yellow), a beta power decrease was not observed between stop-signal onset and SSRT. On the contrary, a beta power decrease was present between stop-signal onset and SSRT when stopping was instructed during movement but did not lead to stopping (red), but we cannot conclude whether this decrease was only related to arrival at the response port, or whether this decrease in addition contained a reactive stopping component, resulting in accumulation of two different beta power decrease processes. The latter scenario could explain why the beta power decrease was strongest in this condition.

Generally, although beta power clearly decreased during stopping, it is important to note that power is always strictly positive, so the decrease should be interpreted as a short but substantial attenuation in beta power relative to the strength of the beta power dynamics prior to trial initiation. Since we contrast power in each frequency band here to its baseline value, the baseline

power across frequency bands is not visible here. It is worthwhile to mention that activity in the beta band is one of the dominant frequencies when animals await the go-signal in both the OFC and STN, similar to the reported dominance of striatal beta power in the study of Leventhal et al. (2012).







Figure 4. Reactive stopping leads to a timed, transient decrease in beta activity in the OFC and STN. Average beta power envelopes for the OFC (**A**) and STN (**B**) are relative to baseline and shown locked to RespT to simplify comparison across conditions, by aligning the expected beta decreases to stopping at port arrival. Conditions are color-matched with outlines of time-frequency windows in Figure 3. Significantly different timepoints between the green and blue conditions are indicated with the blue-green markers on top of the time-axis and represent pFDR-corrected < 0.05 . Color-shaded areas represent standard error of the mean. **A**, Activity in the OFC exhibited a significant decrease in beta power after the stop-signal but before the SSRT when stopping was needed (stop-signal while moving) and successful (green), as compared with when physical stopping was not essential (stop-signal before moving) for making a correct response (blue). **B**, The STN also showed a significant decrease in beta power between the stop-signal and SSRT when stopping was needed and successful (green) as compared with when stopping was instructed but not needed (blue). In contrast to the OFC, beta power did not clearly decrease a second time at port arrival in this condition (green).

Decreased beta phase synchronization between OFC and STN during stopping

Like the OFC and STN beta power reduction, phase synchronization across channels between the OFC and STN in the beta frequency range was decreased after the stop-signal and before the SSRT when stopping was successful, and became strongest around 21 Hz. This decreased beta phase synchronization was also present in the other conditions, but was clearly related to arriving at the lateral port as all six conditions had the strongest decreases in phase

synchronization during or around the response (Figure 5A, time-locked to arrival at the port). The average beta phase synchronization envelope for successful stopping demonstrated that phase synchronization quickly dropped shortly after stop-signal onset. During reactive stopping, the beta phase synchronization envelope dropped to a minimum of 92% of baseline phase synchronization 53 ms before the SSRT. After false-discovery rate correction for multiple comparisons, this decreased beta phase synchronization reached statistical significance around 200 ms after stop-signal onset and lasted until around 67 ms after the SSRT, in comparison to the stopping condition where stopping was not needed and not executed right after the stop-signal, as there was no ongoing movement to terminate (blue). Then, the decrease in beta phase synchronization was shortly released, whereafter beta phase synchronization decreased again reflecting the response-related decrease in beta phase synchronization (Figure 5B).

> **Figure 5.** The OFC and STN transiently decouple in the beta band during reactive stopping. **A**, Time-frequency ISPC showed a response-related decrease in phase synchronization in all six conditions, but only appeared between the stop-signal (STOP) and SSRT when stopping was needed and successful (top right, green). As go trials do not have a stop-signal and SSRT, their theoretical position given the SSD in that session is printed in gray. Other plot elements as in Figures 3 and 4. **B**, Beta ISPC envelopes time-locked to response. When stopping was needed and successful (green), beta phase synchronization significantly decreased as compared with when stopping was instructed but not needed (blue). Significance indicated as in Figure 4. Color-shaded areas represent standard error of the mean. **C**, During stopping, the majority of OFC–OFC and OFC–STN channel pairs changed their connectivity strength relative to baseline (89 and 63%, respectively), while only 19% of STN–STN channel pairs changed their connectivity strength. (Left) Connectivity diagrams show ISPC values relative to baseline between all possible OFC–STN channel pairs (color, suprathreshold connection strength; gray, subthreshold connection, locations correspond to electrode position in the array; dot size, hubness; see Materials and Methods). (Right) Average intra-OFC (right top), intra-STN (left bottom), and interareal (right bottom) suprathreshold connection strength (color, strength; fraction of colored square, proportion of suprathreshold connections).



We visualized beta phase synchronization across animals between all possible channel pairs for successful stopping (green) and the stopping condition where stopping was not needed as the rat did not release from the central port yet (blue), specifically in the time period just before the SSRT when beta phase synchronization decreased significantly relative to baseline. Following the distribution of all channel pair beta connectivities, all absolute beta connectivities higher than a standard deviation above the median (.053) were considered suprathreshold. When stopping was instructed but not needed, only one out of all possible channel pairs had a suprathreshold connection relative to baseline. During successful stopping 63% of all OFC-STN channel pairs had suprathreshold phase synchronization changes relative to baseline with an ΔISPC average of $-.079$, while 89% OFC-OFC channel pairs changed their beta phase synchronization with an ΔISPC average of $-.080$, and 18% of STN-STN channel pairs with an ΔISPC average of $-.067$. Remarkably, none of the suprathreshold channel pairs showed an increase in beta phase synchronization in both conditions. To explore the degree of centrality of all the channels in the network during reactive stopping, hubness for each channel was computed as the proportion of suprathreshold connections a given channel had with all possible channels from the other brain area. Across animals, hubness analysis showed that most OFC channels had comparable amounts of suprathreshold connections with STN channels during reactive stopping, as represented by the comparable dot sizes across most OFC channels (hubness range across OFC channels: 33-83%; almost all $>50\%$). On the other hand, STN channels showed more diversity in the amount of suprathreshold connectivities to other channels (hubness range across STN channels: 0-100%), as 12.5% of STN channels lacked a suprathreshold connection with the OFC, while 25% of STN channels were connected to all possible OFC channels (Figure 5C). The discrepancy between OFC and STN hubness values may suggest that the OFC acts more like a homogenous entity during reactive stopping, because most of the OFC channels had suprathreshold connections with the majority of STN channels at the time of stopping. On the other hand, the STN seems to be more specialized as reflected in a broader range of hubness values: STN channels are either functionally connected with a large majority of OFC channels during reactive stopping, or they are barely involved.

Discussion

In this study we recorded neural activity in the form of local field potentials in the rat OFC and STN during a reactive stopping task, and showed that both areas decrease their ongoing beta power specifically during stopping. This reduction in power began right after stop-signal onset and peaked well before the SSRT, and only occurred when there was ongoing movement. We further showed that beta phase synchronization between the OFC and STN reduced during stopping execution. These insights refine our understanding of the roles played by the OFC and STN in the hyperdirect pathway. While supporting the involvement of these areas in this fast pathway, the results suggest that there is a precisely timed phase desynchronization between OFC and STN during the execution of stopping, contrasting with the previously suggested notion of strengthened synchronization put forth by other papers (Alegre et al., 2013; Aron et al., 2016).

Beta power in OFC and STN is decreased instead of increased during stopping

The decrease in beta power in both the OFC and STN during stopping, consistent across all animals, confirmed that local beta dynamics in the OFC and STN play an important role in terminating ongoing movements. Previous studies that recorded electrophysiology during a stopping task have not been unanimous about the timing and sign of beta activity. Some reported effects after the SSRT (Hubbard & Sahakyan, 2023; Leventhal et al., 2012; Swann et al., 2009; Swann et al., 2011), and some before (Swann et al., 2012; Wagner et al., 2018; Wessel et al., 2016). Some studies showed relative increases in beta power in the STN, but these effects were derived from contrasting two conditions with decreases in beta power (Alegre et al., 2013), or the effects were related to the suppression of response initiation as seen in go/no-go tasks (Kühn et al., 2004; Leventhal et al., 2012). In fact, while Wagner et al. (2018) reported beta power increases over the right frontal cortex, beta power *decreases* were also observed between stop-signal onset and SSRT.

Besides our stop-related beta power decreases, beta power also decreased substantially when the animals were approaching the response port, while beta power only marginally decreased when they started moving towards the go-signal side (Figure 3). This latter, marginal decrease could be explained by animals not going all-in, as they know a stop-signal may be presented while moving to the go-signal side. And, Leventhal et al. (2012) showed that beta

power decreases and increases are not consistently linked to (the initiation of) movement or the absence of movement. The beta power decrease at response port arrival in all conditions suggested to us that this may also be associated with a stopping mechanism.

The OFC and STN show beta phase desynchronization during stopping

Ongoing synchronization was attenuated shortly after stop-signal onset and reached peak desynchronization just before SSRT. Remarkably, this decreased synchronization was observed in most channel-pairs within the OFC and for the majority of channel-pairs in the OFC and STN, while it occurred for only a minority of channel-pairs in the STN. While to our knowledge there are no studies thus far that recorded from the OFC (or inferior frontal cortex in humans) and STN simultaneously while investigating their beta phase synchronization during reactive stopping, the hypothesis of increased beta coherence supporting stopping was plausible given the various studies showing increases in beta power in stopping tasks in frontal cortices and the subthalamic nucleus (Aron et al., 2016; Zavala et al., 2015). Our study confirmed that beta synchronization between the OFC and STN was modulated specifically during stopping, again both during reactive stopping and response port approach. However, ongoing synchronization between the OFC and STN decreased instead of increased during stopping, which did not happen when there was no movement to be stopped.

Beta desynchronization may allow for functionally isolated neuronal activity for stopping

Our results raised the following questions: Is the beta power decrease and the beta phase desynchronization associated with an 'action of stopping'? Or is it opening a window for changing the motor plan? Below we argue that both interpretations are plausible. We speculate that 1) baseline synchronization before stopping prevents the interference of ongoing movement dynamics, and 2) that temporary desynchronization allows for or reflects more functionally isolated neurons in the STN to affect specific routes downstream to engage specific muscles for stopping ongoing movement.

Beta power decreases in sensorimotor areas and the basal ganglia are traditionally associated with movement and muscle contraction changes (Barone & Rossiter, 2021; Kilavik et al., 2013). Given this traditional view, seeing stopping as an active motor response would actually predict beta

power to decrease, because a stopping action requires extensive muscular involvement if there is considerable propulsive force to be stopped (Ellis et al., 2014) and antagonizing muscles are even recruited when movement barely begun (Atsma et al., 2018). The 'active stopping' hypothesis explains why beta power decreases at response port arrival, as the animals need to decelerate to prevent collision with the response port, and is further supported by stronger beta power decreases in OFC and STN during fast go trials as compared to slow go trials. However, the hyperdirect pathway may display distinct beta dynamics as opposed to the traditional view, and Leventhal et al. (2012) showed that beta power dynamics are not consistently linked to (the initiation of) movement or the absence of movement.

In line with the 'status quo' hypothesis of Engel and Fries (2010) and supported by empirical evidence from Leventhal et al. (2012), OFC-STN beta synchronization at baseline may prevent the planned or ongoing movement from being changed or stopped, as maintenance of oscillatory activity in the STN could prevent specific functionally-isolated STN neurons from signaling to the internal globus pallidus to inhibit specific motor actions. We speculate that during maintained local beta synchronization in the STN, neuronal activity is highly correlated and may reflect non-specificity in neuronal firing to downstream areas like the internal globus pallidus. The STN is somatotopically organized (Nambu et al., 1996; Nambu et al., 2002) and STN neurons have been shown to function as uncorrelated parallel processing units (Steiner et al., 2019). So, when the STN oscillates, it is unable to orchestrate the specific complex demands during stopping, because some muscles need inhibition, while other groups of muscles need to contract. However, when OFC-STN synchronization is lifted shortly during stopping, this could cause the STN to desynchronize its local activity, allowing for or reflecting functional isolation of specific neurons to actively signal to the internal globus pallidus, in turn inhibiting thalamic projections to the motor cortex for those muscles that need to stop their contractions. It may be questioned whether decreased synchronization can be associated with increased neuronal firing downstream, but (Courtemanche et al., 2003) showed that increased task-related spike activity in the striatum can occur during disengagement from synchronized beta oscillations, Zavala et al. (2017) showed that decreased beta oscillations in STN were associated with enhanced spike-phase locking, and Lipski et al. (2017) suggested that spike-phase locking in beta can be suppressed in the population of STN neurons as a whole, while a selection of STN neurons can increase spike-phase locking in beta.

Limitations and future directions

It has been shown that the STN receives projections from frontal cortices (Nambu et al., 1996; Nambu et al., 2000; Nambu et al., 2002), and Chen et al. (2020) showed that stop-related field potentials in the frontal cortex preceded those in the STN through a monosynaptic connection. Therefore, we hypothesize that decreased synchronization between OFC and STN causes local desynchronization in the STN, allowing STN neurons to affect specific muscle groups downstream needed for stopping.

The hypothesis that temporary local desynchronization in the STN allows for increased isolated functionality fits well with our data, but the collection of single-neuron activity during a reactive stopping task is needed to verify this hypothesis. Our electrodes were optimized for recording local field potentials and were not suitable for extracting spike data. As beta power does not necessarily correlate positively with striatal neuronal activity (Courtemanche et al., 2003; Lipski et al., 2017; Zavala et al., 2017), it would be worthwhile to study this relationship in a reactive stopping task. Future studies could investigate whether the OFC is causative for establishing local desynchronization in the STN, by using interference methods like optogenetics or electrical stimulation. These methods allow for temporally precise interference, which is needed in a motor circuitry that serves many different roles in a short period of time during motor tasks. Gaining more insights into what happens in downstream areas like the internal globus pallidus and substantia nigra pars reticulata could help to better understand how reductions in beta power in the OFC and STN and desynchronization between them supports downstream signaling needed for stopping. To our knowledge, this is the first study reporting a decrease in beta power in the OFC and STN, along with OFC-STN beta phase desynchronization, specifically occurring before and during the action of stopping.

Chapter 4

General discussion

The overarching goal of this dissertation was to better understand the underlying neural oscillatory mechanisms of reactive stopping in the hyperdirect pathway, by making use of dual-area, multi-electrode recordings in rats performing a multi-session stop-signal task. The first aim was to optimize the rodent stop-signal task design, the second aim was to find out whether within-animal varying stopping speeds were statistically meaningful, and the third, major aim was to establish an electrophysiological rodent model with high temporal and anatomical precision to learn about the oscillatory mechanisms of reactive stopping in the orbitofrontal cortex (OFC) and subthalamic nucleus (STN), two key areas of the hyperdirect pathway. Here, I will summarize the key findings and contributions of this dissertation, discuss open ends and methodological limitations, and present future directions for the field.

Optimizations to the rodent stop-signal task

In the past, several rodent studies presented the stop-signal directly after leaving the central port (Bryden et al., 2012; Bryden & Roesch, 2015), or selected a delay from a set of fixed stop-signal delays (Bari et al., 2009; Mayse et al., 2014). However, as recommended by Verbruggen et al. (2019), the stop-signal delay ideally adapts after every stop trial based on performance to reliably estimate the stop-signal reaction time (SSRT). Therefore, we implemented an adaptive stop-signal delay in our task. In addition, rats could initiate trials by themselves (as long as the inter-trial interval of 3 seconds elapsed) so they were only receiving go- and stop-signals when they were motivated to perform in trials, and go-signal onset was far less predictable due to an added random jitter, as compared to a fixed go-signal onset where rats are able to anticipate the response. As opposed to typical human stop-signal tasks where only the reaction time is recorded (time between go-signal and button press response), in our adapted version of the stop-signal task rats had to hold their nose in the central port for an unpredictable amount of time and had to release after go-signal onset, which allowed us to separate the classical reaction time into release time and response time. This could be useful for future studies investigating reactive stopping in both humans and animals, as the onset of movement (at release time) is very well captured, and allows for dissociating reactive stop-related activity during movement versus more proactive stopping in anticipation of movement, as demonstrated in chapter 3. One way to implement this in a classical computer task for humans is to ask participants to hold the spacebar for a variable amount of time, then

present the go-signal, after which they have to press a key on either the left or right side of the keyboard, where the keys are equally distant to the spacebar. Further, video tracking of the hand movements can be used to get more insight into the detailed behavioral dynamics.

While functional magnetic resonance imaging studies pushed the field forward in knowing which areas are involved in reactive stopping, adding multi-site, multi-electrode recordings in our adapted stop-signal task for rodents improved temporal precision, which allowed us to extract stop-related activity at a time-scale that fitted the speed of the cognitive task better. The custom-designed electrodes were shaped to the anatomical curvature of the OFC and STN, which helped to improve the anatomical precision as opposed to electrophysiological methods (e.g., electroencephalography) that were used before. Altogether, combining the task optimizations as recommended by Verbruggen et al (2019) with an electrophysiological rodent model allowed us to improve anatomical and temporal precision, and has given us insights that were not reported before in the field of reactive stopping.

Summary of experimental chapters

As intracranial electrophysiology was recorded on a daily basis, we acquired many sessions of stop-signal task data. While some studies made use of multi-session stop-signal task data (e.g., see Hall et al., 2022; You et al., 2023; Thunberg et al., 2024), they usually consisted of two within-subject repetitions of the task. However, in this dissertation we acquired many sessions of data (on average 36.4 per animal), and addressed in **chapter 2** whether stopping speed is a trait or state. Little was known about whether stopping speed was a fixed within-animal characteristic or a day-to-day variable state, but numerous studies used the stop-signal task in the past to approximate general inhibitory control capability of individuals (e.g., clinical vs. non-clinical participants, see Lipszyc & Schachar (2010) for a meta-analysis), suggesting it was assumed (consciously or unconsciously) to be a fixed trait within an individual. Under identical experimental conditions, we demonstrated that within-subject SSRTs significantly changed from session to session, and showed that several circumstantial factors like motivation, shared motor dynamics and attention could play a role in these changing stopping speeds. Although it is highly sensible to have a single SSRT estimate when reactive stopping is investigated with a single session, it turns out to be better to treat reactive stopping speed

as a state-like characteristic that is far from fixed within an individual. Armed with the knowledge that SSRTs are far from trait-like in identical experimental conditions, we analyzed our stop-signal task data in **chapter 3** with this in mind; we computed all variables, including SSRT, on a session-by-session basis.

Despite the fact that several studies and review articles hinted towards increases in oscillatory beta activity in the OFC and STN, and increased synchronization between them during reactive stopping, we observed (to our surprise) a *decrease* instead of an *increase* in local oscillatory beta activity in the OFC and STN, specifically after the stop-signal and before the estimated average SSRT. Not only local oscillatory beta activity decreased its power, but ongoing interareal synchronization in the beta range between the OFC and STN *decreased* right after stop-signal onset and before stopping execution. This stop-related decrease in local beta power and interareal synchronization only emerged when stopping was needed, which is, when the stop-signal was presented during ongoing movement, as compared to when the animal was not yet moving towards the go-signal side. This study shows that the hypothesis of increased hyperdirect pathway activity during reactive stopping may need revision, as we did not even observe a single channel pair between the OFC and STN that *increased* interareal synchronization during successful stopping. We speculated that temporary desynchronization between the OFC and STN reflects (or allows for) a situation where local STN neurons become more functionally isolated as opposed to when the ongoing motor plan needs to continue, as we observed maintained beta synchronization when stopping was not required. We speculated that when STN neurons are more functionally isolated (as reflected with locally decreased beta power), they are capable of changing specific downstream muscles to contract (or relax) to aid the cancellation of ongoing movement. In other words, a situation in which ongoing motor plans do not need to be changed may be maintained by beta synchronization, whereas temporary beta desynchronization may help to release the ongoing motor plan and opens up the possibility for specific downstream muscle changes to support stopping behavior.

Our speculated hypothesis about functionally isolated STN neurons fits with what is thought to be the working mechanism behind treatment in Parkinson's disease. When patients with Parkinson's disease receive levodopa or deep brain stimulation, pathologically high beta activity is reduced and symptoms are diminished (Salenius et al., 2002; Kühn et al., 2006; Brown, 2007). Our thought is that this reduction in pathologically elevated beta activity recovers

the ability for the hyperdirect pathway to temporally desynchronize local STN activity to affect specific muscles downstream for stopping, as the STN is not *stuck* anymore in pathological beta synchronization. It is then capable of increasing and decreasing beta synchronization. As our data suggested, temporarily reduced local beta power and interareal desynchronization is associated with reactive stopping. If this is causally related to stopping, it makes sense that stopping is much more difficult when beta synchronization is constantly elevated.

Altogether, we can conclude from this dissertation that the field might need revision on two stopping themes: 1) stopping speed is a state-like, and not trait-like, characteristic that meaningfully changes from time to time, and 2) rather than *increased* beta-band activity in the hyperdirect pathway, reactive stopping seems to be facilitated by *decreased* local beta-band activity and interareal synchronization.

The complexity resides in the details

As pointed out in the general introduction, previous studies investigating reactive stopping have shown inconsistencies. We believe that some inconsistencies may have arisen from 1) different analytical approaches to extract stop-related activity from stop-signal task data, and from 2) erroneously assigning stop-related activities as *contributors* to stopping behavior, while the timing of those neural signatures does not support this. As already pointed out before by Isherwood et al. (2023), there are inconsistencies in the reactive stopping literature. Here, we attempt to provide possible explanations for these inconsistencies and hope that they serve as entries for a better understanding of reactive stopping.

While we were discussing our experimental task design and behavioral data analyses, we realized that simply comparing successful stop trials with unsuccessful stop trials would be problematic for two reasons. First, in unsuccessful stop trials, the stop-signal triggered the stop process too late relative to the go process to enable the animal to successfully stop in time, but that does not mean stop-related activity was absent while the animal was moving towards the (incorrect) lateral port. When contrasting successful stopping with this condition, this would mean that stop-related activity was subtracted from stop-related activity. Second, in unsuccessful stop trials, the

time-window containing stop-related activity overlapped with activity related to arrival at the lateral port. This overlap was of course not surprising because the stop process was triggered too late relative to the go process, causing the animal to finalize the go movement. However, contrasting successful stopping with this condition would mean subtracting arrival-related activity that was not present at that moment in the successful stopping trial, as that arrival at the lateral port happens later in successful stop trials. This would introduce relative activity that is unrelated to stopping. In fact, if we would have contrasted successful stopping with unsuccessful stopping (green vs. red in chapter 3), we would observe an *increase* in beta power in the OFC and STN. But, as argued above, this contrast is problematic and therefore difficult to interpret.

Another camp of researchers argues that successful stopping should be contrasted with slow go trials, and unsuccessful stopping with fast go trials. This makes sense, as one wants to get rid of go-related activity in the stop trial to be left with stop-related activity. The idea behind pairing successful stop trials with slow go trials, while pairing unsuccessful stop trials with fast go trials, comes from the fact that successful stop trials were successful because the go process was slow enough relative to the stop process, and unsuccessful stop trials were unsuccessful because the go process was too fast relative to the stop process. Thus, by contrasting with the matching go process speed, one gets rid of go-related activity optimally. However, when we inspected the timing of different events in these two conditions, we realized it would make much more sense to contrast with stop trials where response port arrival and SSRT were better matched, because this allowed us to properly separate arrival-related activity from stop-related activity. After all, each trial contains go-related activity, so each contrast will remove go-related activity as a consequence.

Lastly, we noticed that some studies assigned changes in beta activity to stopping, while we argue that this is unlikely when the significant change in activity is happening *after* the SSRT. We hypothesize that post-stopping activity reflects stopping-adjacent operations such as error monitoring, reward anticipation, or learning. Another possibility may be that previously reported increases in beta-band activity after SSRT are reflections of a subcortical post-movement beta rebound (Leventhal et al., 2012; Kilavik et al., 2013; Schmidt et al., 2019). In our data we also observed that beta power returned back to higher (baseline) levels of beta power after stopping, and

after arrival at the lateral port. However, we only considered activity to be possible contributors to reactive stopping if it happened between the stop-signal and before the SSRT.

Methodological limitations and future directions

While our stop-signal task was optimized for reliably extracting SSRTs, it was not optimized for assessing *why* SSRTs display state-like characteristics. Therefore, our analyses in chapter 2 had to remain rather correlative. Future studies are required to obtain more conclusive explanations for the mechanisms that underlie varying within-subject stopping speeds. As discussed in chapter 2, the field would benefit from study designs where cognitive states are experimentally controlled and manipulated, for example by comparing groups with different motivational levels (such as thirst in rats) or different levels of attention (by manipulating length of sleep in humans). We cautiously concluded that higher levels of SSRT estimate reliabilities were co-occurring with lower levels of go trial response time variabilities, higher levels of go trial response time skewness and stop trial accuracies of at least 50%, but we are not sure whether the first two factors are easy to manipulate by researchers. Achieving 50% accuracy on stop trials however, is possible when using an adaptive stop-signal delay with enough trials included in the session. While the majority of sessions in chapter 2 had accuracies around 50%, a minority had stop trial accuracies more distant from 50%. This tends to happen in sessions with fewer trials, resulting in incomplete stop-signal delay titrations, in turn causing stop trial accuracies to deviate from 50%. We speculate rats did fewer trials when they were rather quickly satisfied with the amount of water rewards they received (remember trials were initiated by rats themselves), or when they were bored with the task or tired. For future studies, we recommend attempting to increase the number of trials per session by decreasing reward volumes, using wireless recording systems to prevent fatigue as much as possible, or incorporating other strategies to keep rats engaged in the task for longer. An adaptive stop-signal delay in a task with plenty of trials has been proven to be advantageous for SSRT reliability (Verbruggen et al., 2019), so including enough trials in a stop-signal task that adapts difficulty based on stop trial performance is the least one can do to boost SSRT estimate reliability.



Unfortunately our custom-made multi-electrode recordings were not optimized for recording single-neuron activity such as spikes. In chapter 3 we speculated that local STN desynchronization may be the mechanism for stopping behavior, but this can only be confirmed with spike recordings in the STN during stop-signal task execution. Ideally combined with confirmation about which muscles are affected by those recorded neurons, it could be elucidated whether our hypothesis holds or not. As a follow up, it would be worthwhile investigating whether local beta desynchronization (e.g., reduced beta power) in the OFC is causative for local beta desynchronization in the STN. While some animal studies used optogenetics to learn more about the role of the frontal and striatal brain areas in motor behavior in general (Kravitz et al., 2010; Burguière et al., 2013; Guillaumin et al., 2021; Yoon et al., 2014; Yoon et al., 2016)), studies did not include the stop-signal task specifically to find out their causative role in reactive stopping. We think the field would benefit from interference methods like optogenetics (Rossi et al., 2015), where excitability of the OFC-to-STN projection can be inhibited, during which local spike-activity in the STN can be recorded. If it happens to be the case that the OFC is causative for local desynchronization at the level of the STN, and local STN desynchronization supports reactive stopping, this could be verified in the stop-signal task with well-timed interference during task execution, for example by lowering the excitability of the projection to the STN specifically when stopping is requested. Additionally, it would be very insightful to also know whether reactive stopping is still effective when OFC-STN desynchronization can be prevented with excitatory optogenetic stimulation, as this will tell whether interareal desynchronization plays a causal role in reactive stopping. It may be possible that general inhibition of the OFC-to-STN projection is not specific enough when one part of the STN needs to affect downstream muscles differently than another part of the STN due its somatotopic organization (Nambu et al., 1996; Nambu et al., 2002) and uncorrelated parallel processing units (Steiner et al., 2019). In this case, there is demand for a more sophisticated interference method that does not act like a big hammer on local and interareal synchronization.

On a more general note it would be insightful to include video recordings of animals or humans performing the stop-signal task. The advent of easy-to-use, markerless video tracking tools such as DeepLabCut or SLEAP, bring this analysis into close reach on existing data, for example our data (https://data.ru.nl/collections/di/dcmn/DSC_000422.jth_mc_034, doi:10.34973/9na6-fp67) that includes video data from rats performing the stop-signal task. Especially

like in our stop-signal task, animals have to stop their initiated action to one side, and turn to the other side to get a reward. Video data could act as a sanity check for SSRT estimates, as we expect that the ongoing motion should be fully stopped at SSRT, or right before SSRT. In addition, when researchers time-lock physiological data to single-session SSRTs (which we did not do as we time-locked to stop-signal and response at the lateral port for reasons explained in chapter 3), we anticipate that the physiological signature (for example beta power) occurs in a shorter time-window as compared to when an across-session SSRT would be used, as results in chapter 2 indicated that stopping speeds meaningfully change within individuals, which implies that an across-session SSRT estimate is not well-matched with the true stopping speed in individual sessions.

Concluding remarks

When I started working on this project, the literature strongly suggested that one would observe an *increase* in local beta power in the OFC and STN and an *increase* in interareal beta synchronization between those two regions associated with reactive stopping. To my surprise, it turned out to be different, and needed quite some time to trust my data, although I was 100% confident about how the data were acquired and processed. After all, I hope two main outcomes will endure in the minds of reactive stopping researchers: 1) stopping speed is a state, so treat it as such when studying reactive stopping, and 2) reactive stopping may *not* be implemented through increases in beta power and synchronization, but rather by decreases. I believe my supervisors Mike and Bernhard gave me many stop-signals in the meantime, so at some point it is time for me to conclude. I hope other researchers feel encouraged and motivated to further determine what really goes on in stopping minds, and that this dissertation may be a helpful starting point for further in-depth investigation.

Appendices

References

Nederlandse samenvatting

Research data management

About the author

Portfolio

List of publications

Donders Graduate School for
Cognitive Neuroscience

References

- Albin RL, Young AB, Penney JB (1989) The functional anatomy of disorders of the basal ganglia. *Trends Neurosci* 18:63–64. doi:10.1016/0166-2236(89)90074-x
- Alegre M, Lopez-Azcarate J, Obeso I, Wilkinson L, Rodriguez-Oroz MC, Valencia M, Garcia-Garcia D, Guridi J, Artieda J, Jahanshahi M, Obeso JA (2013). The subthalamic nucleus is involved in successful inhibition in the stop-signal task: a local field potential study in Parkinson's disease. *Experimental Neurology* 239:1–12. doi:10.1016/j.expneurol.2012.08.027
- Aron AR, Behrens TE, Smith S, Frank MJ, Poldrack RA (2007) Triangulating a cognitive control network using diffusion-weighted magnetic resonance imaging (MRI) and functional MRI. *J Neurosci* 27:3743–3752. doi:10.1523/JNEUROSCI.0519-07.2007
- Aron AR, Herz DM, Brown P, Forstmann BU, Zaghoul K (2016) Frontosubthalamic circuits for control of action and cognition. *J Neurosci* 36:11489–11495. doi:10.1523/JNEUROSCI.2348-16.2016
- Aron AR, Poldrack RA (2006) Cortical and subcortical contributions to stop signal response inhibition: role of the subthalamic nucleus. *J Neurosci* 26:2424–2433. doi:10.1523/JNEUROSCI.4682-05.2006
- Aron AR, Robbins TW, Poldrack RA (2004) Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences* 8:170–177. doi:10.1016/j.tics.2004.02.010
- Aron AR, Robbins TW, Poldrack RA (2014) Inhibition and the right inferior frontal cortex: one decade on. *Trends Cogn Sci (Regul Ed)* 18:177–185. doi:10.1016/j.tics.2013.12.003
- Atsma J, Maij F, Gu C, Medendorp WP, Corneil BD (2018). Active braking of whole-arm reaching movements provides single-trial neuromuscular measures of movement cancellation. *The Journal of Neuroscience* 38:4367–4382. doi:10.1523/JNEUROSCI.1745-17.2018
- Bari A, Eagle DM, Mar AC, Robinson ESJ, Robbins TW (2009) Dissociable effects of noradrenaline, dopamine, and serotonin uptake blockade on stop task performance in rats. *Psychopharmacology (Berl)* 205:273–283. doi:10.1007/s00213-009-1537-0
- Bari A, Mar AC, Theobald DE, Elands SA, Oganya KCNA, Eagle DM, Robbins TW (2011) Prefrontal and monoaminergic contributions to stop-signal task performance in rats. *J Neurosci* 31:9254–9263. doi:10.1523/JNEUROSCI.1543-11.2011
- Barone J, Rossiter HE (2021). Understanding the role of sensorimotor beta oscillations. *Frontiers in Systems Neuroscience* 15:1–7. doi:10.3389/fnsys.2021.655886
- Brown P (2007) Abnormal oscillatory synchronisation in the motor system leads to impaired movement. *Curr Opin Neurobiol* 17:656–664. doi:10.1016/j.conb.2007.12.001
- Bryden DW, Burton AC, Kashtelyan V, Barnett BR, Roesch MR (2012) Response inhibition signals and miscoding of direction in dorsomedial striatum. *Front Integr Neurosci* 6:69. doi:10.3389/fnint.2012.00069
- Bryden DW, Roesch MR (2015) Executive control signals in orbitofrontal cortex during response inhibition. *J Neurosci* 35:3903–3914. doi:10.1523/JNEUROSCI.3587-14.2015
- Burguière E, Monteiro P, Feng G, Graybiel AM (2013) Optogenetic stimulation of lateral orbitofronto-striatal pathway suppresses compulsive behaviors. *Science* 340:1243–1246. doi:10.1126/science.1232380
- Carlson JS, Jensen CM, Widaman KF (1983) Reaction time, intelligence, and attention. *Intelligence* 7:329–344. doi:10.1016/0160-2896(83)90008-9

- Chen W, de Hemptinne C, Miller AM, Leibbrand M, Little SJ, Lim DA, Larson PS, Starr PA (2020) Prefrontal-subthalamic hyperdirect pathway modulates movement inhibition in humans. *Neuron* 106:579–588. doi:10.1016/j.neuron.2020.02.012
- Cohen MX (2019). A better way to define and describe Morlet wavelets for time-frequency analysis. *Neuroimage* 199:81–86. doi:10.1016/j.neuroimage.2019.05.048
- Courtemanche R, Fujii N, Graybiel AM (2003). Synchronous, focally modulated β band oscillations characterize local field potential activity in the striatum of awake behaving monkeys. *The Journal of Neuroscience*, 23:11741–1752. doi:10.1523/JNEUROSCI.23-37-11741.2003
- Coxon JP, Van Impe A, Wenderoth N, Swinnen SP (2012) Aging and inhibitory control of action: cortico-subthalamic connection strength predicts stopping performance. *J Neurosci* 32:8401–8412. doi:10.1523/JNEUROSCI.6360-11.2012
- Curley LB, Newman E, Thompson WK, Brown TT, Hagler DJ, Akshoomoff N, Reuter C, Dale AM, Jernigan TL (2018) Cortical morphology of the pars opercularis and its relationship to motor-inhibitory performance in a longitudinal, developing cohort. *Brain Struct Funct* 223:211–220. doi:10.1007/s00429-017-1480-5
- Delorme A, Makeig S (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods* 134:9–21. doi:10.1016/j.jneumeth.2003.10.009
- Doekemeijer RA, Dewulf A, Verbruggen F, Boehler CN (2023) Proactively Adjusting Stopping: Response Inhibition is Faster when Stopping Occurs Frequently. *J Cogn* 6:22. doi:10.5334/joc.264
- Eagle DM, Baunez C, Hutcheson DM, Lehmann O, Shah AP, Robbins TW (2008) Stop -signal reaction-time task performance: role of prefrontal cortex and subthalamic nucleus. *Cereb Cortex* 18:178–188. doi: 10.1093/cercor/bhm044
- Eagle DM, Baunez C (2010) Is there an inhibitory-response-control system in the rat? Evidence from anatomical and pharmacological studies of behavioral inhibition. *Neurosci Biobehav Rev* 34:50–72. doi:10.1016/j.neubiorev.2009.07.003
- Ellis RG, Sumner BJ, Kram R (2014). Muscle contributions to propulsion and braking during walking and running: insight from external force perturbations. *Gait & Posture* 40:594–599. doi:10.1016/j.gaitpost.2014.07.002
- Engel AK, Fries P (2010). Beta-band oscillations - signalling the status quo? *Current Opinion in Neurobiology* 20:156–165. doi:10.1016/j.conb.2010.02.015
- Feola TW, Wit H de, Richards JB (2000) Effects of d-Amphetamine and ethanol on a measure of behavioral inhibition in rats. *Behav Neurosci* 114:838–848. doi:10.1037/0735-7044.114.4.838
- Fife KH, Gutierrez-Reed NA, Zell V, Bailly J, Lewis CM, Aron AR, Hnasko TS (2017) Causal role for the subthalamic nucleus in interrupting behavior. *eLife* 6. doi:10.7554/eLife.27689
- Forstmann BU, Keuken MC, Jahfari S, Bazin P-L, Neumann J, Schäfer A, Anwender A, Turner R (2012) Cortico-subthalamic white matter tract strength predicts interindividual efficacy in stopping a motor response. *Neuroimage* 60:370–375. doi:10.1016/j.neuroimage.2011.12.044
- França ASC, van Hulten JA, Cohen MX (2020). Low-cost and versatile electrodes for extracellular chronic recordings in rodents. *Heliyon* 6:1–6. doi:10.1016/j.heliyon.2020.e04867
- Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn Sci (Regul Ed)* 9:474–480. doi:10.1016/j.tics.2005.08.011

- Fries P (2015) Rhythms for Cognition: Communication through Coherence. *Neuron* 88:220–235. doi:10.1016/j.neuron.2015.09.034
- Gerfen CR, Surmeier DJ (2011) Modulation of striatal projection systems by dopamine. *Annu Rev Neurosci* 34:441–466. doi:10.1146/annurev-neuro-061010-113641
- Gordi VM, Drueke B, Gauggel S, Antons S, Loevenich R, Mols P, Boecker M (2019) Stopping Speed in the Stop-Change Task: Experimental Design Matters! *Front Psychol* 10:279. doi:10.3389/fpsyg.2019.00279
- Guillaumin A, Serra GP, Georges F, Wallén-Mackenzie Å (2021) Experimental investigation into the role of the subthalamic nucleus (STN) in motor control using optogenetics in mice. *Brain Res* 1755:147226. doi:10.1016/j.brainres.2020.147226
- Hall A, Jenkinson N, MacDonald HJ (2022) Exploring stop signal reaction time over two sessions of the anticipatory response inhibition task. *Exp Brain Res* 240:3061–3072. doi:10.1007/s00221-022-06480-x
- Haynes WIA, Haber SN (2013) The organization of prefrontal-subthalamic inputs in primates provides an anatomical substrate for both functional specificity and integration: implications for basal ganglia models and deep brain stimulation. *J Neurosci* 33:4804–4814. doi:10.1523/JNEUROSCI.4674-12.2013
- Hubbard RJ, Sahakyan L (2023) Differential recruitment of inhibitory control processes by directed forgetting and thought substitution. *J Neurosci* 43:1963–1975. doi:10.1523/JNEUROSCI.0696-22.2023
- Isherwood SJS, Kemp S, Miletić S, Stevenson N, Bazin PL, Forstmann BU (2023) The canonical stopping network: Revisiting the role of the subcortex in response inhibition. *eLife* 12:1–31. doi:10.7554/eLife.88652.1
- Jahanshahi M, Obeso I, Rothwell JC, Obeso JA (2015) A fronto-striato-subthalamic-pallidal network for goal-directed and habitual inhibition. *Nature Rev Neuroscience* 16:719–732. doi:10.1038/nrn4038
- Kilavik BE, Zaepffel M, Brovelli A, MacKay WA, Riehle A (2013) The ups and downs of β oscillations in sensorimotor cortex. *Exp Neurol* 245:15–26. doi:10.1016/j.expneurol.2012.09.014
- Kravitz AV, Freeze BS, Parker PRL, Kay K, Thwin MT, Deisseroth K, Kreitzer AC (2010) Regulation of parkinsonian motor behaviours by optogenetic control of basal ganglia circuitry. *Nature* 466:622–626. doi:10.1038/nature09159
- Kühn AA, Kupsch A, Schneider G-H, Brown P (2006) Reduction in subthalamic 8–35 Hz oscillatory activity correlates with clinical improvement in Parkinson's disease. *Eur J Neurosci* 23:1956–1960. doi:10.1111/j.1460-9568.2006.04717.x
- Kuhn AA, Williams D, Kupsch A, Limousin P, Hariz M, Schneider GH, Yarrow K, Brown P (2004) Event-related beta desynchronization in human subthalamic nucleus correlates with motor performance. *Brain* 127:735–746. doi:10.1093/brain/awh106
- Leventhal DK, Gage GJ, Schmidt R, Pettibone JR, Case AC, Berke JD (2012) Basal ganglia beta oscillations accompany cue utilization. *Neuron* 73:523–536. doi:10.1016/j.neuron.2011.11.032
- Lipski WJ, Wozny TA, Alhourani A, Kondylis ED, Turner RS, Crammond DJ, Richardson RM (2017). Dynamics of human subthalamic neuron phase-locking to motor and sensory cortical oscillations during movement. *J Neurophysiol* 118:1472–1487. doi:10.1152/jn.00964.2016
- Lipszyc J, Schachar R (2010) Inhibitory control and psychopathology: a meta-analysis of studies using the stop signal task. *J Int Neuropsychol Soc* 16:1064–1076. doi:10.1017/S1355617710000895

- Logan GD, Cowan WB (1984) On the ability to inhibit thought and action: A theory of an act of control. *Psychol Rev* 91:295–327. doi:10.1037/0033-295X.91.3.295
- Madsen KS, Johansen LB, Thompson WK, Siebner HR, Jernigan TL, Baaré WFC (2020) Maturational trajectories of white matter microstructure underlying the right presupplementary motor area reflect individual improvements in motor response cancellation in children and adolescents. *Neuroimage* 220:117105. doi:10.1016/j.neuroimage.2020.117105
- Magill PJ, Sharott A, Bevan MD, Brown P, Bolam JP (2004) Synchronous unit activity and local field potentials evoked in the subthalamic nucleus by cortical stimulation. *J Neurophysiol* 92:700–714. doi:10.1016/j.neuroimage.2020.117105
- Malvea A, Babaei F, Boulay C, Sachs A, Park J (2022) Deep brain stimulation for Parkinson's Disease: A Review and Future Outlook. *Biomed Eng Lett* 12:303–316. doi:10.1007/s13534-022-00226-y
- Mayse JD, Nelson GM, Park P, Gallagher M, Lin S-C (2014) Proactive and reactive inhibitory control in rats. *Front Neurosci* 8:104. doi:10.3389/fnins.2014.00104
- Nambu A, Takada M, Inase M, Tokuno H (1996) Dual somatotopical representations in the primate subthalamic nucleus: evidence for ordered but reversed body-map transformations from the primary motor cortex and the supplementary motor area. *J Neurosci* 16:2671–2683. doi:10.1523/JNEUROSCI.16-08-02671.1996
- Nambu A, Tokuno H, Hamada I, Kita H, Imanishi M, Akazawa T, Ikeuchi Y, Hasegawa N (2000). Excitatory cortical inputs to pallidal neurons via the subthalamic nucleus in the monkey. *J Neurophysiol* 84:289–300. doi:10.1152/jn.2000.84.1.289
- Nambu A, Tokuno H, Takada M (2002) Functional significance of the cortico-subthalamo-pallidal “hyperdirect” pathway. *Neurosci Res* 43:111–117. doi:10.1016/S0168-0102(02)00027-5
- Neubert F-X, Mars RB, Buch ER, Olivier E, Rushworth MFS (2010) Cortical and subcortical interactions during action reprogramming and their related white matter pathways. *Proc Natl Acad Sci USA* 107:13240–13245. doi:10.1073/pnas.1000674107
- Parent A, Hazrati LN (1995) Functional anatomy of the basal ganglia. II. The place of subthalamic nucleus and external pallidum in basal ganglia circuitry. *Brain Res Brain Res Rev* 20:128–154. doi:10.1016/0165-0173(94)00008-D
- Paxinos G, Watson C (2013). *The rat brain in stereotaxic coordinates*, 7th edn. San Diego: Academic Press.
- Prinzmetal W, McCool C, Park S (2005) Attention: reaction time and accuracy reveal different mechanisms. *J Exp Psychol Gen* 134:73–92. doi:10.1037/0096-3445.134.1.73
- Rossi MA, Calakos N, Yin HH (2015) Spotlight on movement disorders: What optogenetics has to offer. *Mov Disord* 30:624–631. doi:10.1002/mds.26184
- Salenius S, Avikainen S, Kaakkola S, Hari R, Brown P (2002) Defective cortical drive to muscle in Parkinson's disease and its improvement with levodopa. *Brain* 125:491–500. doi:10.1093/brain/awf042
- Schaum M, Pinzuti E, Sebastian A, Lieb K, Fries P, Mobascher A, Jung P, Wibral M, Tüscher O (2021) Right inferior frontal gyrus implements motor inhibitory control via beta-band oscillations in humans. *eLife* 10. doi:10.7554/eLife.61679
- Schmidt R, Herrojo Ruiz M, Kilavik BE, Lundqvist M, Starr PA, Aron AR (2019) Beta oscillations in working memory, executive control of movement and thought, and sensorimotor function. *J Neurosci* 39:8231–8238. doi:10.1523/JNEUROSCI.1163-19.2019
- Singh KD (2012) Which “neural activity” do you mean? fMRI, MEG, oscillations and neurotransmitters. *Neuroimage* 62:1121–1130. doi:10.1016/j.neuroimage.2012.01.028

- Steiner LA, Barreda Tomás FJ, Planert H, Alle H, Vida I, Geiger JRP (2019) Connectivity and dynamics underlying synaptic control of the subthalamic nucleus. *J Neurosci* 39:2470–2481. doi:10.1523/JNEUROSCI.1642-18.2019
- Surmeier DJ, Ding J, Day M, Wang Z, Shen W (2007) D1 and D2 dopamine-receptor modulation of striatal glutamatergic signaling in striatal medium spiny neurons. *Trends Neurosci* 30:228–235. doi:10.1016/j.tins.2007.03.008
- Swann N, Poizner H, Houser M, Gould S, Greenhouse I, Cai W, Strunk J, George J, Aron AR (2011) Deep brain stimulation of the subthalamic nucleus alters the cortical profile of response inhibition in the beta frequency band: a scalp EEG study in Parkinson's disease. *J Neurosci* 31:5721–5729. doi:10.1523/JNEUROSCI.6135-10.2011
- Swann N, Tandon N, Canolty R, Ellmore TM, McEvoy LK, Dreyer S, DiSano M, Aron AR (2009) Intracranial EEG reveals a time- and frequency-specific role for the right inferior frontal gyrus and primary motor cortex in stopping initiated responses. *J Neurosci* 29:12675–12685. doi:10.1523/JNEUROSCI.6135-10.2011
- Swann NC, Cai W, Conner CR, Pieters TA, Claffey MP, George JS, Aron AR, Tandon N (2012) Roles for the pre-supplementary motor area and the right inferior frontal gyrus in stopping action: electrophysiological responses and functional and structural connectivity. *Neuroimage* 59:2860–2870. doi:10.1016/j.neuroimage.2011.09.049
- Thomsen BLC, Jensen SR, Clausen A, Karlsborg M, Jespersen B, Løkkegaard A (2020) Deep brain stimulation in parkinson's disease: still effective after more than 8 years. *Mov Disord Clin Pract* 7:788–796. doi:10.1002/mdc3.13040
- Thunberg C, Wiker T, Bundt C, Huster RJ (2024) On the (un)reliability of common behavioral and electrophysiological measures from the stop signal task: Measures of inhibition lack stability over time. *Cortex* 175:81–105. doi:10.1016/j.cortex.2024.02.008
- Verbruggen F, Chambers CD, Logan GD (2013) Fictitious inhibitory differences: how skewness and slowing distort the estimation of stopping latencies. *Psychol Sci* 24:352–362. doi:10.1177/0956797612457390
- Verbruggen F, Logan GD (2009) Models of response inhibition in the stop-signal and stop-hange paradigms. *Neurosci Biobehav Rev* 33:647–661. doi:10.1016/j.neubiorev.2008.08.014
- Verbruggen F et al. (2019) A consensus guide to capturing the ability to inhibit actions and impulsive behaviors in the stop-signal task. *eLife* 8:1–26. doi:10.7554/eLife.46323
- Wagner J, Wessel JR, Ghahremani A, Aron AR (2018) Establishing a right frontal beta signature for stopping action in scalp EEG: implications for testing inhibitory control in other task contexts. *J Cogn Neurosci* 30:107–118. doi:10.1162/jocn_a_01183
- Weber S, Salomoni SE, St George RJ, Hinder MR (2024) Stopping speed in response to auditory and visual stop signals depends on go signal modality. *J Cogn Neurosci* 36:1395–1411. doi:10.1162/jocn_a_02171
- Wessel JR, Ghahremani A, Udupa K, Saha U, Kalia SK, Hodaie M, Lozano AM, Aron AR, Chen R (2016) Stop-related subthalamic beta activity indexes global motor suppression in Parkinson's disease. *Mov Disord* 31:1846–1853. doi:10.1002/mds.26732
- Williams BR, Ponesse JS, Schachar RJ, Logan GD, Tannock R (1999) Development of inhibitory control across the life span. *Dev Psychol* 35:205–213. doi:10.1037//0012-1649.35.1.205
- Yoon HH, Min J, Hwang E, Lee CJ, Suh J-KF, Hwang O, Jeon SR (2016) Optogenetic Inhibition of the Subthalamic Nucleus Reduces Levodopa-Induced Dyskinesias in a Rat Model of Parkinson's Disease. *Stereotact Funct Neurosurg* 94:41–53. doi:10.1159/000442891

- Yoon HH, Park JH, Kim YH, Min J, Hwang E, Lee CJ, Suh J-KF, Hwang O, Jeon SR (2014) Optogenetic inactivation of the subthalamic nucleus improves forelimb akinesia in a rat model of Parkinson disease. *Neurosurgery* 74:533-540; discussion 540. doi:10.1227/NEU.0000000000000297
- You Y, Failla A, van der Kamp J (2023) No training effects of top-down controlled response inhibition by practicing on the stop-signal task. *Acta Psychol (Amst)* 235:103878. doi:10.1016/j.actpsy.2023.103878
- Yu C, Cassar IR, Sambangi J, Grill WM (2020) Frequency-Specific Optogenetic Deep Brain Stimulation of Subthalamic Nucleus Improves Parkinsonian Motor Behaviors. *J Neurosci* 40:4323-4334. doi:10.1523/JNEUROSCI.3071-19.2020
- Zavala B, Zaghloul K, Brown P (2015) The subthalamic nucleus, oscillations, and conflict. *Movement Disorders*, 30:328-338. doi:10.1002/mds.26072

Nederlandse samenvatting

Reactief stoppen

Stel je voor, je bent in de buitenlucht aan het hardlopen en je nadert een kruispunt met verkeerslichten. Het verkeerslicht staat op groen, dus alles in je hersenen is erop voorbereid om het kruispunt over te steken. Terwijl je nog maar een paar meter verwijderd bent van het kruispunt springt het verkeerslicht op rood. Nu zullen je hersenen in een hele korte tijd al je spieren moeten instrueren om te stoppen met voortbewegen om te voorkomen dat je in botsing komt met een andere weggebruiker. Deze situatie wordt *reactief stoppen* genoemd, omdat je geïnstrueerd wordt om te stoppen met voortbewegen als reactie op het zien van het rode verkeerslicht, het *stop-sig-naal*. Reactief stoppen wordt moeilijker des te later het stop-sig-naal wordt gepresenteerd. Er wordt niet zozeer meer van je spieren gevraagd, maar hoe dichter je bij het kruispunt bent als het verkeerslicht op rood springt, des te minder tijd er over is om nog op tijd te stoppen. Met andere woorden, het moment van het rode licht in combinatie met je nabijheid tot het kruispunt is bepalend in hoe waarschijnlijk het is dat je nog op tijd kunt stoppen. Natuurlijk is het een stuk makkelijker om nog op tijd te stoppen als je je ver van het kruispunt bevindt als het verkeerslicht op rood springt.

Alhoewel reactief stoppen weinig lijkt te vragen van onze mentale capaciteit, gebeurt er wel degelijk veel in de hersenen. Allereerst moet het stop-sig-naal waargenomen worden en moet het geassocieerd worden met de context waarin het zich begeeft. Zo is een rood licht in de disco geen instructie om te stoppen met dansen, terwijl een rood verkeerslicht direct geassocieerd wordt met het stoppen met rijden, hardlopen of wandelen. Nadat de hersenen hebben geconstateerd dat stoppen gewenst is, moet dit voorbereid worden in hersenen. In tegenstelling tot *proactief stoppen*, waarbij de hersenen continu anticiperen en voorbereid zijn op stoppen, vraagt reactief stoppen van de hersenen heel snel een onvoorbereide reactie op een niet-geanticipeerde situatie. Dit suggereert dat reactief stoppen zo geïmplementeerd moet zijn in de hersenen dat het de juiste spieren aanstuurt terwijl je je houding en balans weet te behouden, en zonder dat het stop-proces zó lang duurt dat je pas tot stilstand komt óp het kruispunt.

Hoe wordt reactief stoppen bestudeerd in een lab?

Reactief stoppen wordt al enige decennia onderzocht door hersenwetenschappers, waarbij specifiek gebruik wordt gemaakt van een computertaak. In

computertaken worden proefpersonen (deelnemers aan een experiment) vaak gevraagd om een knop op een toetsenbord in te drukken als ze een bepaalde stimulus zien, de respons. Een stimulus is een waarneembare prikkel, bijvoorbeeld een pieptoon of een lichtje dat gaat branden. Een dergelijke stimulus-respons taak wordt heel vaak herhaald, en deze herhalingen worden *trials* genoemd. Reactief stoppen wordt meestal onderzocht met de stop-sigitaal taak. Proefpersonen worden dan gevraagd om op een knop te drukken als ze een *go-sigitaal* zien op een scherm, vaak een visuele stimulus in de vorm van een lichtje dat gaat branden. Vergelijkbaar met het verkeerslicht is dat dus een groen verkeerslicht. Het maakt eigenlijk niet uit of dit een visuele stimulus is of bijvoorbeeld een geluidssigitaal, zolang het maar geassocieerd wordt met het maken van een meetbare reactie, zoals het indrukken van een toets op een toetsenbord. Deze trials worden *go trials* genoemd. Bij een klein deel van de trials (meestal 25%), wordt het go-sigitaal gevolgd door een andere stimulus, het stop-sigitaal (bijvoorbeeld een rood licht), dat de proefpersoon instrueert om de geplande en geïnitieerde respons op het go-sigitaal af te breken. Kortom, het stop-sigitaal vraagt de proefpersoon om de knop niet in te drukken, ook al werd dit initieel wel geïnstrueerd met het go-sigitaal. Deze trials worden *stop trials* genoemd. Zoals bij het voorbeeld met het verkeerslicht bepaalt de mate van vertraging tussen het go-sigitaal en het stop-sigitaal de moeilijkheid. Als het stop-sigitaal vrij laat wordt gepresenteerd ten opzichte van het go-sigitaal, dan is het relatief moeilijk om nog te stoppen, omdat de respons op het go-sigitaal al bijna is voltooid. Andersom wordt het makkelijker als de vertraging tussen het go- en stop-sigitaal verkleind wordt, want de beweging richting de knop voor een respons is nog maar net op gang gekomen. In dat geval is het makkelijker om op tijd te stoppen met bewegen richting de knop. De vertraging tussen het go-sigitaal en het stop-sigitaal noemen we de *stop-sigitaal delay*, of stop-sigitaal vertraging. In de stop-sigitaal taak is deze vertraging idealiter niet vaststaand, maar verandert deze iedere keer na een stop trial. Als de proefpersoon erin slaagt te stoppen en dus niet de knop in te drukken, wordt bij de volgende stop trial de stop-sigitaal vertraging vergroot, waardoor het stop-sigitaal later wordt gegeven. Hierdoor wordt het moeilijker nog op tijd te stoppen. Andersom wordt na een niet-succesvolle stop trial de vertraging verkleind bij de volgende stop trial, waardoor het stop-sigitaal vroeger wordt gepresenteerd. Daardoor wordt het makkelijker om op tijd te stoppen. Op deze manier zal elke proefpersoon de stop-sigitaal taak als even moeilijk ervaren, omdat de stop-sigitaal vertraging op den duur zal gaan settelen rond een waarde die zorgt voor 50% correcte responsen op stop trials: de proefpersoon heeft evenveel kans om te slagen in stoppen als dat het kans

heeft niet te slagen in stoppen. Een ander voordeel hiervan is dat het moment van het stop-signaal enigszins onvoorspelbaar wordt, en daarmee dus niet geanticipeerd kan worden.

Maar hoe weten we nou hoe snel een proefpersoon is in reactief stoppen? Een correcte 'respons' is namelijk de afwezigheid van een respons, omdat het stop-signaal immers instrueert om de knop niet in te drukken. Kortom, de snelheid van stoppen (*stop-signal reaction time*, stop-signaal reactietijd, SSRT) is onmeetbaar want er wordt per definitie geen knop ingedrukt. Wetenschappers hebben een wiskundig model ontwikkeld die het mogelijk maakt de snelheid van stoppen nauwkeurig in te schatten op basis van hoe snel de proefpersoon reageert op go trials, en de kans dat de proefpersoon de knop indrukt terwijl er een stop-signaal werd gepresenteerd. In essentie wordt er gezocht naar de reactietijd die overeenkomt met de 'interne' stop-respons. De SSRT kan geschat worden door: 1) alle reactietijden op go trials te ordenen van snel naar langzaam, 2) de reactietijd in die ordening te vinden die overeenkomt met de kans dat de proefpersoon de knop indrukt terwijl er een stop-signaal werd gepresenteerd, 3) de stop-signaal vertraging af te trekken van deze reactietijd. Stel je voor, een proefpersoon slaagde erin om in 55% van de stop trials te stoppen. Dan was de kans dat de proefpersoon de knop indrukte bij een stop trial 45% ($100\% - 55\% = 45\%$). De go trial reactietijd die overeenkomt met die kans in een sessie van 200 go trials is de 90ste snelste go trial reactietijd. Vervolgens kan de SSRT berekend worden door de gemiddelde stop-signaal vertraging van de 90ste snelste go trial reactietijd af te trekken; het stop-proces kan immers pas starten als het stop-signaal is gepresenteerd.

Wat hebben wetenschappers al ontdekt over reactief stoppen?

Wetenschappers kunnen de stop-signaal taak combineren met hersenmetingen. Zo werd MRI (*magnetic resonance imaging*, magnetische resonantie beeldvorming) in het verleden gebruikt om te kijken welke hersengebieden betrokken zijn bij reactief stoppen. Zo kwam men tot de ontdekking dat de orbitofrontale cortex en de subthalamische nucleus actief zijn tijdens reactief stoppen, alsmede dat sterkere activiteit samenhangt met het vermogen sneller te kunnen stoppen (kortom, een lagere stop-signaal reactietijd). De orbitofrontale cortex is een hersengebied dat zich vrijwel direct boven de oogkassen bevindt, terwijl de subthalamische nucleus een heel klein hersengebied is dat zich diep en midden in de hersenen bevindt, ongeveer ter hoogte van de bovenste aanhechting van de oren. Andere onderzoekers lieten zien dat beschadigingen in de orbitofrontale cortex en subthalamische nucleus

bij ratten reactief stoppen bemoeilijkt, wat suggereert dat deze hersengebieden belangrijk zijn voor reactief stoppen. Ook zijn er aanwijzingen dat de integriteit van de hersenbaan (een soort verbindingsweg) tussen de orbitofrontale cortex en subthalamische nucleus samenhangt met hoe goed men kan stoppen; als de hersenbaan van betere kwaliteit is (zoals beter asfalt en meer rijstroken op een snelweg), is men ook beter in stoppen in een stop-signaal taak.

Naast MRI hebben hersenwetenschappers ook EEG (elektro-encefalografie) en ECoG (elektro-corticografie) ingezet om reactief stoppen beter te begrijpen. Dat zijn technieken waarbij zeer zwakke elektrische stroompjes gemeten kunnen worden door heel gevoelige sensoren aan de buitenkant van de schedel of hersenen. Deze zeer zwakke stroompjes worden gegenereerd door hersencellen (daar communiceren ze namelijk mee). Men kwam tot de ontdekking dat hersengolven met een specifieke frequentie meer worden waargenomen tijdens reactief stoppen, namelijk bètagolven. Deze specifieke hersengolven, voortkomend uit heel veel zwakke stroompjes van heel veel hersencellen bij elkaar, hebben een frequentie van 12 tot 30 hertz. Dat wil zeggen dat de hersengolf 12 tot 30 keer op en neer gaat in een seconde (zoals stembanden in het strottenhoofd ook met een bepaalde frequentie heen en weer kunnen trillen, resulterend in verschillende toonhoogtes). Deze bètagolven werden ook in sterkere mate waargenomen in succesvolle stop trials ten opzichte van stop trials waarbij het niet lukte te stoppen. Bij patiënten met Parkinson die diepe hersenstimulatie (*deep brain stimulation*, DBS) krijgen in de subthalamische nucleus als onderdeel van een behandeling, kwam men erachter dat bètagolven in de orbitofrontale cortex ook sterker aanwezig zijn rond het moment van stoppen, alsmede dat ze sneller waren in stoppen, in vergelijking met wanneer de hersenstimulatie uit werd gezet.

Deze observaties suggereerden dat bètagolven in de orbitofrontale cortex en subthalamische nucleus belangrijk zijn voor reactief stoppen. Hieruit voortkomend kwam de theorie tot stand dat de orbitofrontale cortex waarschijnlijk het stoppen initieert door de subthalamische nucleus een signaal te geven door middel van bètagolven, waardoor de orbitofrontale cortex en subthalamische nucleus samen gaan synchroniseren in de bèta-frequentie. Synchronisatie wil in deze context zeggen dat de bètagolven samen op en neer gaan in zowel de orbitofrontale cortex en subthalamische nucleus. Het idee heerst in de hersenwetenschappen dat dit communicatie tussen hersengebieden kan bewerkstelligen. Door synchronisatie in de bèta-frequentie tussen de orbitofrontale cortex en subthalamische nucleus zou

de subthalamische nucleus het stoppen in gang kunnen zetten door andere hersengebieden de opdracht te geven de spieren dusdanig te activeren en deactiveren dat het stoppen in werking wordt gezet.

Wat wordt er nog niet begrepen over reactief stoppen?

Alhoewel voorgaande wetenschappelijke studies suggereren dat versterkte bètagolven een belangrijke rol spelen in reactief stoppen, er is niet zo veel bewijs op het gebied van elektrofysiologie in zowel de orbitofrontale cortex en subthalamische nucleus, specifiek na het stop-signaal en voor het daadwerkelijke moment van stoppen (de stop-signaal reactietijd). Sommige studies observeerden veranderingen in bètagolven vóór het stopmoment, terwijl andere studies juist veranderingen in bètagolven observeerden ná het stopmoment. Zo waren er ook studies die niet tegelijkertijd in de orbitofrontale cortex en subthalamische nucleus maten, of ontbrak het aan anatomische of temporele precisie. Anatomische precisie gaat over hoe zeker men kon zijn van de bron van de hersengolven, terwijl de temporele (of tijds-)precisie gaat over hoe zeker men kon zijn dat de activiteit die zij maten ook plaatsvond op het moment van stoppen. Soms waren wetenschappelijke resultaten gebaseerd op weinig proefpersonen, of werden veranderingen in hersengolven op groepsniveau gedreven door maar een deel van de proefpersonen. Bovendien, al zouden beide hersengebieden sterkere bètagolven hebben tijdens reactief stoppen, dat wil nog niet zeggen dat ze ook met elkaar communiceren (en synchroniseren) door middel van bètagolven. Ook betekent het niet automatisch dat bètagolven in de subthalamische nucleus een direct gevolg zijn van bètagolven in de orbitofrontale cortex.

Daarom hebben we ratten getraind op het uitvoeren van een stop-signaal taak. Rattenhersen en mensenhersen hebben heel veel gemeen, en zijn daardoor geschikt voor fundamenteel onderzoek naar hoe de hersenen werken. In tegenstelling tot bij mensen kun je bij ratten relatief makkelijk in de hersenen meten, wat de anatomische precisie ten goede komt, in plaats van aan de oppervlakte van de schedel of hersenen. We hebben elektrodes ontworpen en gemaakt die geschikt zijn om hele zwakke stroompjes in de hersenen te meten, die specifiek de activiteit van de orbitofrontale cortex en subthalamische nucleus kunnen observeren doordat ze ontworpen zijn aan de hand van de vorm en grootte van de betreffende hersengebieden. Nadat de ratten getraind waren op de stop-signaal taak hebben we de zelfgebouwde elektroden geïmplantéerd en waren we in staat om de elektrofysiologie van de hersenen te meten terwijl ze de stop-signaal taak uitvoerden.

Het ontwikkelen van een dergelijk diermodel, zoals dat heet in de vakliteratuur, stelt onderzoekers ook in staat om meer geavanceerde technologie te gebruiken om causale verbanden te onderzoeken. In dierproeven is het namelijk mogelijk interventies te gebruiken die hersenonderzoekers de mogelijkheid biedt om bijvoorbeeld een heel specifiek hersengebied tijdelijk uit te schakelen met optogenetische of elektrische stimulatie. Zo zou kunnen worden aangetoond of bètagolven in de orbitofrontale cortex causaal verband houden met bètagolven in de subthalamische nucleus, door te observeren of bètagolven in de subthalamische nucleus nog wel plaatsvinden tijdens reactief stoppen als de orbitofrontale cortex tijdelijk is uitgeschakeld. Maar, voor de wetenschap zo ver is, moeten we er eerst in slagen tegelijkertijd te meten in de orbitofrontale cortex en subthalamische nucleus van ratten terwijl ze de stop-sigitaal taak kunnen uitvoeren.

De stop-sigitaal taak voor ratten

In tegenstelling tot de stop-sigitaal taak voor humane proefpersonen, waarbij op knoppen kan worden gedrukt, vraagt een stop-sigitaal taak bij ratten om een andere benadering. Allereerst is het onmogelijk ratten uit te leggen wat de bedoeling is, ten tweede zijn ratten niet bepaald gewend om knoppen van een toetsenbord in te drukken terwijl ze naar een beeldscherm kijken. Daarom hebben we een versie gemaakt van de stop-sigitaal taak die kan worden aangeleerd bij ratten in een operante kamer, ook wel een Skinner-box genoemd. De rat bevindt zich dan in een ruimte waarbij door middel van beloningen de rat stap voor stap leert wat het gewenste gedrag is bij het waarnemen van bepaalde lichtjes. Deze ruimte bevat één wand met drie kleine poortjes waar de rat zijn snuit in kan steken. In de stop-sigitaal taak initieert de rat een trial door zijn snuit in de middelste poort te steken voor ongeveer één seconde, waarna een go-sigitaal wordt gegeven door middel van een lichtje dat kort brandt in de linker of rechter poort. De rat moet dan uit de middelste poort komen en zijn snuit in de betreffende poort steken om daar een beloning te krijgen in de vorm van een waterdruppel. De rat is gemotiveerd om deze waterdruppel te krijgen omdat de ratten een deel van hun dagelijkse waterinname krijgen in de stop-sigitaal taak en daardoor dorstig zijn voor aanvang van de taak. In 25% van de trials wordt het go-sigitaal gevolgd door een stop-sigitaal na een variabele stop-sigitaal vertraging. In dat geval krijgt de rat dus eerst een go-sigitaal in een van de poorten (links of rechts), en ergens onderweg naar de betreffende poort zal het lichtje van de andere poort gaan branden. Dit instrueert de rat om te stoppen met bewegen richting de poort van het eerste lichtje, om vervolgens zijn snuit in de andere poort te

steken voor het krijgen van een beloning. De rat wordt dus gevraagd de initiële beweging te stoppen als reactie op het stop-signaal. Omdat je ratten geen instructies kunt geven zoals humane proefpersonen, wordt de taak stap voor stap aangeleerd door middel van operante conditionering. Eerst leren ze dat het loont om hun snuit in de middelste poort te steken, vervolgens leren ze dat ze hun snuit in de poort moeten steken waar het lichtje brandt, en als laatste stap leren ze dat ze hun beweging naar het eerste lichtje moeten afbreken als het andere lichtje gaat branden. Het duurt zes tot acht weken voordat een rat deze taak kan uitvoeren, waarbij ze iedere werkdag een uur oefenen en in het weekend rust krijgen.

De snelheid van stoppen varieert van moment tot moment

Doordat we zwakke stroompjes maten in de hersenen en elektrofysiologische data veel ruis bevatten, hadden we veel data nodig. Daarom lieten we elke rat heel veel sessies van de stop-signaal taak doen, zodat we ook heel vaak maten wat er gebeurde in de hersenen als de rat werd geïnstrueerd om een ingezette beweging te stoppen. Door de gemiddelde hersenactiviteit van al deze stop trials te berekenen konden we er zo achter komen wat er daadwerkelijk gebeurde in de hersenen, omdat ruis uit het gemiddelde signaal verdwijnt door het willekeurige karakter van ruis. Als gevolg hiervan hadden de ratten heel veel sessies van de stop-signaal taak gedaan, en hadden we dus een geschatte stop-signaal reactietijd per sessie per rat. Daarom vroegen we ons in hoofdstuk 2 van dit proefschrift af in hoeverre de snelheid van reactief stoppen een vaststaande eigenschap is binnen proefpersonen- of dieren, of een eigenschap die onderhevig is aan verandering. Dit was namelijk nog niet of nauwelijks onderzocht door andere wetenschappers, en er werd in de wetenschap tot dusver (bewust of onbewust) vanuit gegaan dat ieder individu een bepaalde stopsnelheid heeft die niet kan veranderen onder gelijke experimentele omstandigheden. We onderzochten of een afzonderlijk geschatte SSRT per sessie statistische meerwaarde had ten opzichte van een enkele geschatte SSRT over alle sessies heen. Resultaten lieten zien dat binnen elk dier de stopsnelheden noemenswaardig verschilden van sessie tot sessie, en de stopsnelheden lieten geen logische trend zien over tijd door bijvoorbeeld het beter worden in de stop-signaal taak. Ook bleek dat minder betrouwbare stopsnelheden samenhangen met meer variabiliteit in de reactietijd van go trials, minder scheefheid in de go trial responstijd distributie, en mindere prestaties in stop trials. Daarnaast zochten we uit welke factoren zouden kunnen bijdragen aan veranderende stopsnelheden, en hebben laten zien dat motivatie, gedeelde bewegingsdynamieken, en aandacht een rol zouden kunnen spelen.

Op basis van deze resultaten adviseren we onderzoekers om de snelheid van stoppen te benaderen als een eigenschap die aan verandering onderhevig is, omdat onze data overtuigend lieten zien dat stopsnelheden alles behalve vaststaand zijn in gelijke experimentele omstandigheden. Als onderzoekers in toekomstig onderzoek in meerdere sessies hersenactiviteit meten tijdens de stop-signaal taak, zullen zij in staat zijn op een preciezere wijze stop-gerelateerde activiteit uit hersenactiviteit te halen, omdat een sessie-specifieke benadering van de stopsnelheid rekening houdt met de verandering waaraan het onderhevig is, terwijl de klassieke benadering van stopsnelheden dat niet doet.

Bètagolven in en tussen de orbitofrontale cortex en subthalamische nucleus worden minder sterk tijdens reactief stoppen

Zoals eerder beschreven waren er weinig studies die zowel met anatomische als temporele precisie tegelijkertijd hebben gemeten in de orbitofrontale cortex en subthalamische nucleus, en bleken veranderingen in bètagolven niet altijd plaats te vinden tijdens het stoppen, maar na het stoppen. Ook is synchronisatie tussen de beide gebieden nog niet bestudeerd tijdens reactief stoppen. In hoofdstuk 3 van dit proefschrift onderzochten we de bètagolven in de orbitofrontale cortex en subthalamische nucleus, en of er synchronisatie in de bèta-frequentie plaatsvindt tussen de orbitofrontale cortex en subthalamische nucleus tijdens reactief stoppen. We vonden dat zowel de orbitofrontale cortex en subthalamische nucleus bijna gelijktijdig een verminderde sterkte van de bètagolven lieten zien tijdens het reactief stoppen, ten opzichte van het moment dat de ratten nog aan het wachten waren op het go-signaal, en ten opzichte van het moment dat ze nog aan het bewegen waren richting de kant waar het go-signaal werd gepresenteerd. Met andere woorden: in tegenstelling tot een versterking van bètagolven vonden we dat deze bètagolven juist zwakker werden tijdens reactief stoppen. Ook vonden we dat synchronisatie in de bèta-frequentie verminderde tijdens reactief stoppen. Deze effecten vonden specifiek plaats tussen het stop-signaal en de stop-signaal reactietijd, wat suggereert dat ze daadwerkelijk hebben kunnen bijdragen aan het stopgedrag van de ratten. Voor zover bekend is dit de eerste keer dat er wordt vastgesteld dat lokale bètagolven in de orbitofrontale cortex en subthalamische nucleus en bèta-synchronisatie tussen deze gebieden in sterkte verminderen tijdens reactief stoppen.

We speculeren dat aanhoudende verhoogde bètagolven er voor zorgen dat alle neuronen (hersencellen) in de subthalamische nucleus het bewegingsplan in stand houden zoals die op dat moment is, zoals we observeerden als reactief

stoppen niet benodigd is. Anderzijds denken we dat een tijdelijke vermindering van bètagolven in de subthalamische nucleus, zoals bij reactief stoppen, ruimte biedt voor neuronen in de subthalamische nucleus om specifieke spieren aan te spannen en te ontspannen om reactief stoppen mogelijk te maken. Anders gezegd vermoeden we dat neuronen in de subthalamische nucleus synchrone bèta-activiteit laten zien als beweging op dat moment niet onderbroken hoeft te worden, maar dat deze synchrone bèta-activiteit kortstondig wordt opgeheven doordat alle neuronen ieder op een unieke manier actief zijn als gevolg van de verminderde bètagolven uit de orbitofrontale cortex, met als gevolg dat neuronen in de subthalamische nucleus minder synchrone activiteit vertonen. Het zou dus kunnen dat de orbitofrontale cortex een cruciale rol speelt bij het in toom en synchroon houden van neurale activiteit in de subthalamische nucleus, maar dat deze synchronisatie kort wordt losgelaten om reactief stoppen te bewerkstelligen. Echter, deze mogelijke causale relatie kan alleen worden aangetoond met interventietechnieken zoals optogenetische of elektrische stimulatie, waarbij men uitzoekt of een tijdelijke verstoring van activiteit in de orbitofrontale cortex door deze stimulatie direct gevolg heeft voor activiteit in de subthalamische nucleus, en of dit gevolgen heeft voor het vermogen reactief te stoppen.

Research data management

Ethics

All animal procedures in chapter 2 and 3 were approved by the Animal Welfare Body of the Radboud University Nijmegen and the Animal Experiment Committee (CCD No. AVD10300 2016 482, Project No. 2015-0129), according to national and international laws, to protect welfare under experimental conditions.

FAIR principles

Unprocessed electrophysiological data, video data and behavioral data from chapter 2 and 3 are openly available at the Radboud Data Repository of Radboud University Nijmegen (https://data.ru.nl/collections/di/dcmn/DSC_000422.jth_mc_034, doi:10.34973/9na6-fp67). Please follow the instructions that can be found at the helppage from the Radboud Data Repository (<https://data.ru.nl/doc/help/helppages/visitor-manual/vm-request-access.html>).

MATLAB code to clean the electrophysiological data is provided, as well as example code to inspect video data. The code used to run the stop-signal task in the Skinner box is supplied too, in case one is interested in how the data is acquired. To use the included MATLAB code, a MATLAB license is needed (<https://nl.mathworks.com/products/matlab.html>, although it is possible to look at the code in free software like Notepad). Some parts of the included code depend on the EEGLAB toolbox (<https://eeglab.org>, freely available for download).

Data will be preserved for 10 years and is intentionally uploaded in raw unprocessed format so researchers are not limited by any means to treat the data as they would like. All provided code is commented and described in the READMEFIRST.txt file, which is added to the stored data collection.

About the author

Jordi ter Horst pursued his bachelor's degree in psychology at the Radboud University in Nijmegen. For his bachelor thesis, he worked on developing a tactile brain-computer interface under supervision of prof. dr. Peter Desain. Next, he was accepted into the research master's programme in cognitive neuroscience at the Radboud University in Nijmegen. In his master thesis, he investigated conflict-related theta activity in the midfrontal cortex in a rat model of response conflict, under supervision of dr. Nils Zuiderveen Borgesius and dr. Mike X. Cohen, which was nominated for the Radboud University thesis award. After obtaining his master's degree *cum laude* in 2019, Jordi was granted the Donders TopTalent Grant to set up his own PhD research at Radboudumc to investigate the neural oscillatory mechanisms underlying reactive stopping. Under supervision of dr. Mike X. Cohen and dr. Bernhard Englitz, he finalized his PhD dissertation in 2024. Besides doing research at the Donders Institute and the department of cognitive neuroscience at Radboudumc, Jordi enjoyed teaching neuroanatomy, signal processing, neurophysiology, and data analysis during his bachelor's, master's and PhD at Radboud University. Jordi will pursue his career outside academia.

Portfolio

COURSES & WORKSHOPS	Organizer	Date
Laboratory Animal Science	Radboudumc	November 2018
Scientific Integrity Course	Donders Graduate School	February 2020
Graduate School Introduction Day	Donders Graduate School	April 2020
Education in a Nutshell	Radboud University	June 2020
Graduate School Day I	Donders Graduate School	July 2020
Internship supervision for CNS	Radboudumc	December 2020
Docentprofessionalisering: feedback geven	Radboudumc	April 2021
Analyzing Neural Time Series	Radboud University	July 2021
Graduate School Day II	Donders Graduate School	June 2022
HPC/Linux workshop	Radboud University	November 2022

TRAINING	Supervisor
Coding and analyzing in MATLAB	Mike X Cohen
Operant conditioning	Nils Zuiderveen Borgesius
Electrode design and manufacturing	Sjef van Hulten, Arthur de França
Surgical implantation	Sjef van Hulten, Arthur de França, Paul Anderson
Brain extraction	Sjef van Hulten
Histology and microscopy	Sjef van Hulten, Bram van Geenen

CERTIFICATES	Organization
Art. 9 (ID: 058/17_97_2018)	FELASA

AWARDS	Organization
Donders TopTalent Grant (€240.000)	Radboudumc/Donders Institute

CONFERENCES	Contribution
Donders Discussions, RU	Oral presentation
Joint Deep Brain Meeting, UvA	Poster presentation
Dutch Neuroscience Meeting	Poster presentation
Dutch Neuroscience Meeting	Poster presentation
TNU Afternoon, Radboudumc	Oral and poster presentation

COMMITTEES	Date
Donders PhD Council - PhD representative	2019 - 2020
Translational Neuroscience Unit - advisory board	2020 - 2023

TEACHING	Role
Neuroanatomy (SOW-DGCN08)	Organizing & teaching
Cell Biophysics (NWI-MOL137)	Teaching
Translational Neuroscience (MED-MIN16)	Teaching & supervision
Recording and analyzing physiological signals (MED-B1RES1)	Teaching
Measuring and modelling reflexes (MED-B2RS3)	Teaching
Master's internship (2021, Mauricio Diaz-Ortiz Jr.)	Supervision
Master's internship (2022, Kasif Işık)	Supervision
Master's internship (2023, Marie-Anna Sedlinská)	Supervision



List of publications

Ter Horst, J., Boillot, M., Cohen, M. X., & Englitz, B. (2024). Decreased beta power and OFC-STN phase synchronization during reactive stopping in freely behaving rats. *The Journal of Neuroscience*, 44, 1-13. doi:10.1523/JNEUROSCI.0463-24.2024

Boillot, M., **ter Horst, J.**, López, J. R., Di Fazio, I., Steens, I. L. M., Cohen, M. X., & Homberg, J. R. (2024). Serotonin transporter knockout in rats reduces beta- and gamma-band functional connectivity between the orbitofrontal cortex and amygdala during auditory discrimination. *Cerebral Cortex*, 34, 1-18. doi:10.1093/cercor/bhae334

Preprint:

Ter Horst, J., Cohen, M. X., & Englitz, B. (2024). Stopping speed as state, not trait: Exploring within-animal varying stopping speeds in a multi-session stop-signal task. *bioRxiv*. doi:10.1101/2024.09.05.611370

Donders Graduate School for Cognitive Neuroscience

For a successful research institute, it is vital to train the next generation of scientists. To achieve this goal, the Donders Institute for Brain, Cognition and Behaviour established the Donders Graduate School in 2009. The mission of the Donders Graduate School is to guide our graduates to become skilled academics who are equipped for a wide range of professions. To achieve this, we do our utmost to ensure that our PhD candidates receive support and supervision of the highest quality.

Since 2009, the Donders Graduate School has grown into a vibrant community of highly talented national and international PhD candidates, with over 500 PhD candidates enrolled. Their backgrounds cover a wide range of disciplines, from physics to psychology, medicine to psycholinguistics, and biology to artificial intelligence. Similarly, their interdisciplinary research covers genetic, molecular, and cellular processes at one end and computational, system-level neuroscience with cognitive and behavioral analysis at the other end. We ask all PhD candidates within the Donders Graduate School to publish their PhD thesis in the Donders Thesis Series. This series currently includes over 600 PhD theses from our PhD graduates and thereby provides a comprehensive overview of the diverse types of research performed at the Donders Institute. A complete overview of the Donders Thesis Series can be found on our website: <https://www.ru.nl/donders/donders-series>

The Donders Graduate School tracks the careers of our PhD graduates carefully. In general, the PhD graduates end up at high-quality positions in different sectors, for a complete overview see <https://www.ru.nl/donders/destination-our-former-phd>. A large proportion of our PhD alumni continue in academia (>50%). Most of them first work as a postdoc before growing into more senior research positions. They work at top institutes worldwide, such as University of Oxford, University of Cambridge, Stanford University, Princeton University, UCL London, MPI Leipzig, Karolinska Institute, UC Berkeley, EPFL Lausanne, and many others. In addition, a large group of PhD graduates continue in clinical positions, sometimes combining it with academic research. Clinical positions can be divided into medical doctors, for instance, in genetics, geriatrics, psychiatry, or neurology, and in psychologists, for instance as healthcare psychologist, clinical neuropsychologist, or clinical psychologist. Furthermore, there are PhD graduates who continue to work as researchers outside academia, for instance at non-profit or government organizations,

or in pharmaceutical companies. There are also PhD graduates who work in education, such as teachers in high school, or as lecturers in higher education. Others continue in a wide range of positions, such as policy advisors, project managers, consultants, data scientists, web- or software developers, business owners, regulatory affairs specialists, engineers, managers, or IT architects. As such, the career paths of Donders PhD graduates span a broad range of sectors and professions, but the common factor is that they almost all have become successful professionals.

For more information on the Donders Graduate School, as well as past and upcoming defences please visit: <http://www.ru.nl/donders/graduate-school/phd/>

