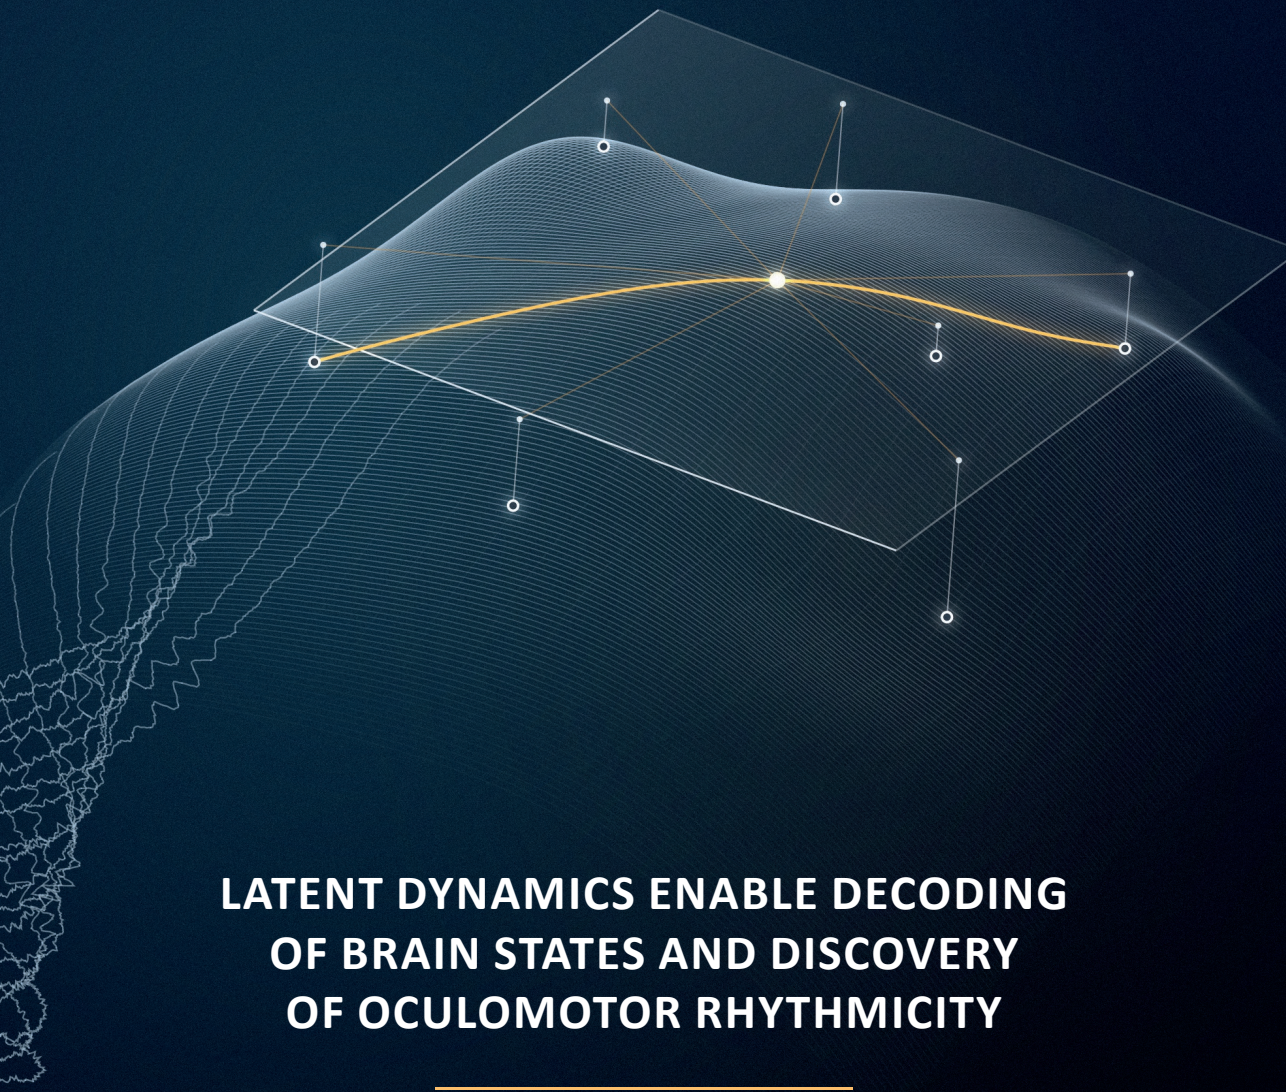

LATENT DYNAMICS ENABLE DECODING OF BRAIN STATES AND DISCOVERY OF OCULOMOTOR RHYTHMICITY

Tim Näher

DONDERS SERIES

RADBOUD UNIVERSITY PRESS

Latent Dynamics Enable Decoding of Brain States and Discovery of Oculomotor Rhythmicity

Tim Volker Näher

Tim Volker Näher

Latent Dynamics Enable Decoding of Brain States and Discovery of Oculomotor Rhythmicity

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: ProefschriftOntwerp.nl
Cover image: Tim Volker Näher
Production coordination & prepress: ProefschriftOntwerp.nl
Printing: ProefschriftMaken/Pumbo

ISBN: 9789465153414
DOI: 10.54195/9789465153414
Free download at: https://doi.org/10.54195/9789465153414

© 2026 Tim Volker Näher

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/.

Latent Dynamics Enable Decoding of Brain States and Discovery of Oculomotor Rhythmicity

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 december 2026
om 16.30 uur precies

door

Tim Volker Näher
geboren op 17 januari 1993
te Neunkirchen (Duitsland)

Promotoren:

Prof. dr. P. Fries

Prof. dr. P.H.M. de Weerd (Universiteit Maastricht)

Manuscriptcommissie:

Prof. dr. W.P. Medendorp

Prof. dr. O. Jensen (University of Oxford, Verenigd Koninkrijk)

Prof. dr. M.W. Tangermann

Latent Dynamics Enable Decoding of Brain States and Discovery of Oculomotor Rhythmicity

Dissertation to obtain the degree of doctor
from Radboud University Nijmegen
on the authority of the Rector Magnificus prof. dr. J.M. Sanders,
according to the decision of the Doctorate Board
to be defended in public on
Monday, December 7, 2026
at 4.30 pm

by

Tim Volker Näher
born on January 17, 1993
in Neunkirchen (Germany)

Supervisors:
Prof. dr. P. Fries
Prof. dr. P.H.M. de Weerd (Maastricht University)

Manuscript Committee:
Prof. dr. W.P. Medendorp
Prof. dr. O. Jensen (University of Oxford, United Kingdom)
Prof. dr. M.W. Tangermann

for Lisa, the love of my life

Contents

Samenvatting		11
Acknowledgements		13
Chapter 1:	General Introduction	15
Chapter 2:	Riemannian Geometry Boosts Functional Near-Infrared Spectroscopy-Based Brain-State Classification Accuracy	29
Chapter 3:	Sensitive and Specific fNIRS-Based Approach for Awareness Detection in Disorders of Consciousness	59
Chapter 4:	Primate Saccade Rhythmicity	81
Chapter 5:	FEF Inhibition Increases Saccade Rhythmicity	113
Chapter 6:	General Discussion	139
	Bibliography	157
Appendix A:	Data Management	175
Appendix B:	Curriculum Vitae	179
Appendix C:	List of Publications	181
Appendix D:	AI Disclosure	183
	Donders Graduate School	185

Samenvatting

Dit proefschrift onderzoekt hoe latente neuronale en gedragsmatige dynamiek kan worden afgeleid uit indirecte, beperkte metingen binnen twee verschillende domeinen van de systeemneurowetenschap. In beide domeinen is het centrale probleem dat het meest in het oog springende kenmerk van het gemeten signaal niet samenvalt met de structuur waarin het onderzochte fenomeen zich bevindt.

Deel I richt zich op functionele nabij-infraroodspectroscopie (fNIRS), een niet-invasieve techniek die hersenactiviteit afleidt uit veranderingen in de cerebrale hemodynamiek. Traditionele fNIRS-analyses behandelen elk meetkanaal afzonderlijk, terwijl hersentoestanden voornamelijk berusten op de interactie-structuur tussen kanalen. Hoofdstuk 2 introduceert een Riemann-geometrisch classificatiekader waarin covariantiematrices van fNIRS-signalen worden behandeld als elementen op een gekromd oppervlakte. Met deze benadering wordt de nauwkeurigheid in het classificeren van acht hersentoestanden bijna verdubbeld ten opzichte van klassieke, op afzonderlijke kanalen gebaseerde methoden. Hoofdstuk 3 past dit kader toe in een klinische proof-of-principle-studie voor de diagnose van bewustzijnsstoornissen, waarbij gesimuleerde respons- en niet-responstoestanden bij gezonde proefpersonen met een sensitiviteit van 100% en een specificiteit van 89% worden onderscheiden.

Deel II behandelt de temporele organisatie van saccadische oogbewegingen tijdens naturalistisch kijken. Hoofdstuk 4 laat zien dat het ontstaan van saccades uit twee componenten bestaat: een lokaal regelmatig interval-genererend proces en een traag fluctuerende saccade-frequentie. Wanneer inter-saccadische intervallen worden herschaald op basis van de lokale frequentie, blijkt dat de saccadische ritmiciteit monotoon toeneemt met de saccadefrequentie bij mensen, makaken en een marmoset. Hoofdstuk 5 toetst causaal de hypothese dat corticale input deze ritmiciteit onregelmatiger maakt. Continue theta-burststimulatie (cTBS) over het rechter frontale oogveld (FEF) verhoogt de saccadische ritmiciteit zonder de saccadefrequentie te beïnvloeden, terwijl gelijktijdige EEG-veranderingen bevestigen dat de manipulatie de lokale FEF-prikkelbaarheid en de connectiviteit met het dorsale aandachtsnetwerk heeft verminderd.

Samen illustreren beide delen dat vergelijkbare problemen met de afleiding van een onderliggende structuur uit data zich voordoen in niet-gerelateerde domeinen van de systeemneurowetenschap, en dat het identificeren van de juiste latente structuur, of deze nu meetkundig of generatief van aard is, vaak het verschil maakt tussen een schijnbare paradox en een interpreteerbaar resultaat.

Acknowledgements

This thesis in its present form would not exist without the contributions, guidance, and support of many people.

First and foremost, I want to thank my supervisor, Pascal Fries. I could not have asked for a better PI. Pascal fostered an environment where curiosity was not just welcomed but actively encouraged, and he made every experiment and idea feel possible. I have learned an extraordinary amount under his mentorship, and I will carry that with me well beyond my PhD.

I owe a special debt to Peter de Weerd, my mentor during my Master's, who set me on this academic path in the first place. It was Peter who first sparked my excitement about perception and attention, and whose encouragement and support ultimately made this PhD a reality.

I also want to thank my brilliant wife, Lisa Bastian. Lisa is not only my partner in life but was also my closest collaborator throughout my PhD. We developed ideas, built models together, and challenged each other's thinking at every turn. Her scientific talent and unwavering support have shaped both this work and the researcher I have become.

I am also thankful to my collaborators Anna Vorreuther, Friederike Breuer, Maximilian Lueckel, Til Ole Bergmann, Yufeng Zhang, Martina Pandinelli, Bettina Sorger, Rainer Goebel, Ziad M. Hafed, Peter Thier, and Silvia Spadacenta, without whom the work presented in this thesis would not have been possible. Their expertise and generosity with their time enriched these projects immeasurably.

I want to thank my colleagues Yufeng Zhang, Martina Pandinelli, and Rasmus Roese for making the day-to-day life of this PhD not only productive but genuinely enjoyable.

Finally, I thank my parents, Ulrike and Stefan, whose continuous support and encouragement have been a constant throughout my life. They gave me the freedom to follow my own path, trusted my choices even when the destination was unclear, and never once wavered in their belief that I would find my way. For that I am more grateful than I can put into words.



1

Chapter 1

General Introduction

This thesis investigates how latent neuronal and behavioral dynamics can be recovered from indirect and constrained measurements across two domains of systems neuroscience. In Part I of the thesis, I introduce geometry-aware analysis methods for functional near-infrared spectroscopy (fNIRS). I demonstrate that distributed brain states are encoded in the interaction structure of hemodynamic signals rather than in information from individual channels. These states can be decoded using representations that respect the intrinsic geometry of the data. Part II addresses the temporal organization of saccadic eye movements during natural viewing, showing that a rhythmic saccade generator is obscured by slow fluctuations in saccade rate. I disentangle these components both analytically and experimentally, by developing a rescaling approach that corrects for saccade rate and by designing a causal transcranial magnetic stimulation (TMS) intervention that modulates rhythmicity while leaving rate unchanged. Despite the significant differences in modality and domain, both lines of work in this thesis share a common inference challenge: the directly observable signals do not immediately reveal the generative process that produced them. Recovering the latent structure of interest therefore requires analytical frameworks that match the generative process rather than the surface features of the measurement.

Part I: Functional Near-Infrared Spectroscopy as a systems neuroscience tool

fNIRS has emerged as a versatile and non-invasive neuroimaging tool for studying human brain functions, in both experimental and applied settings. By measuring task- and state-dependent changes in cerebral hemodynamics, fNIRS can not only provide indirect access to neuronal population activity but also has the combined advantages of spatial specificity, continuous monitoring, and ecological validity (Klein et al., 2024; Kohl et al., 2020; Pinti et al., 2024; Soekadar et al., 2021; Yücel et al., 2025). These properties have made fNIRS particularly attractive for longitudinal measurements, bedside and clinical applications, as well as experiments that are difficult to realize in the more restrictive imaging environments of functional magnetic resonance imaging (fMRI).

In the field of systems neuroscience, fNIRS covers a conceptual space similar to fMRI: both modalities measure vascular responses that arise from the coupling between neuronal activity and local metabolic demand (Logothetis et al., 2001). Therefore, like fMRI, fNIRS does not provide a direct readout of neuronal firing or oscillatory dynamics but rather captures population-level consequences of coordinated neuronal activity. This is often viewed as a limitation. However, being able to measure population-level activity also allows fNIRS to be a highly suitable tool for neuroscientific investigations at the systems level.

Principles of fNIRS measurement and neurovascular coupling

fNIRS is based on the transmission of near-infrared light through biological tissue and the differential absorption characteristics of oxygenated and deoxygenated hemoglobin (Ferrari & Quaresima, 2012; Jöbsis, 1977). In an fNIRS setup, light emitters and detectors are positioned

on the scalp to form measurement channels. Each measurement channel samples changes in hemoglobin concentration from cortical tissue a few centimeters under the skull surface. Neuronal activity induces localized changes in cerebral blood flow and, consequently, oxygenation, leading to characteristic increases in oxygenated hemoglobin (HbO) and decreases in deoxygenated hemoglobin (HbR), which are measured as a function of time.

The physiological foundation of the fNIRS signal is neurovascular coupling. This coupling is based on a multi-stage process that links synaptic and spiking activity to vascular responses (Attwell et al., 2010; Iadecola, 2017). Converging evidence from fMRI and electrophysiology shows that hemodynamic signals are especially sensitive to neuronal gamma-band activity (Logothetis et al., 2001; Niessing et al., 2005; Scheeringa et al., 2011). Changes in the BOLD (blood-oxygen-level-dependent) signal reflect both gamma oscillations and neuronal firing rates for externally induced conditions. However, when BOLD fluctuations arise from internally generated brain states, the BOLD responses preferentially track gamma-band activity while showing a reduced relationship to neuronal firing rates (Koch et al., 2009; Niessing et al., 2005). This relationship positions gamma oscillations as a crucial marker of internal state dynamics.

Although fNIRS does not directly measure neuronal oscillations such as gamma, it captures the same vascular effect of population-level activity that is also captured by the BOLD signal. This means that fNIRS is particularly well suited for studying internally driven brain states, where distributed gamma-associated activity gives rise to structured hemodynamic patterns. Although the temporal dynamics of neurovascular coupling are approximately 2-3 orders of magnitude slower than those observed in electrophysiological recordings (Logothetis, 2008), they still provide sufficient temporal resolution to support the measurement and classification of task-related brain states, spontaneous fluctuations, and condition-dependent differences. Critically, because these signals reflect coordinated activity across neuronal populations, fNIRS measurements form an inherently relational, spatially distributed signal, in which cross-channel structure carries information about the underlying population state.

Historical data analysis approaches in fNIRS

Early fNIRS studies largely adopted analysis strategies from fMRI, which focused on univariate and channel-wise activation mapping using general linear models or block-averaged contrasts (Plichta et al., 2007; Schroeter et al., 2004; Ye et al., 2009). While such approaches have been highly successful for demonstrating task-related cortical engagement and for validating fNIRS as a functional imaging tool, univariate analyses treat each measurement channel as an independent unit and therefore ignore any relations between channels.

More recent developments in fNIRS systems started to push toward high-density recordings (Eggebrecht et al., 2014) and consequently, multivariate approaches are beginning to gain more traction. However, many machine learning methods often still rely on information extracted

independently from each channel. These methods have been applied to classification and decoding tasks, especially in brain-computer interface (BCI) applications (Hong et al., 2018; Naseer & Hong, 2015). While these methods leverage information across channels at the classifier level, they typically assume that features themselves are embedded in a Euclidean space and do not explicitly model interactions between channels.

In contrast to this, fMRI research has increasingly emphasized functional connectivity and network-level analyses (Biswal et al., 1995; Fox & Raichle, 2007). This is motivated by the understanding that cognitive and behavioral states are often encoded in distributed patterns of coordinated activity (Haxby et al., 2001; Kriegeskorte et al., 2008; Norman et al., 2006). Although this understanding is directly transferable to fNIRS, the shift to connectivity-based pipelines into practical fNIRS analyses has remained comparatively limited. Potential explanations for this could be the limited penetration depth of fNIRS (Strangman et al., 2013) and the inability to capture deep sub-cortical brain regions, or simply the reduced number of measurement units, as fMRI can measure thousands of voxels at the same time.

From a systems-neuroscience point of view, brain states are not defined from, and therefore not localized to, single regions, but rather emerge through structured activity across distributed neuronal populations (Ebitz & Hayden, 2021; Vyas et al., 2020). Hemodynamic signals measured with fNIRS are therefore likely best understood as components of a multivariate system whose interactions reflect these underlying brain states. This view suggests that meaningful information is not just confined to features that are derived from individual channels but is also, and potentially primarily, encoded in the interactions across channels and/or their associated features, that need to be considered when analyzing fNIRS data.

Covariance-based representations of fNIRS data

A note on terminology. In this chapter and in the general discussion (Chapter 6), "covariance matrix" is used as an umbrella term for closely related symmetric positive definite (SPD) representations of channel interactions. This includes covariances computed with alternative estimators, correlation matrices, and non-linear kernel matrices such as Laplacian kernels. These matrices share the same underlying geometry and are therefore grouped under a single conceptual label. Chapter 2 uses the reverse convention: there, "kernel matrix" is the umbrella term, and the covariance matrix is treated as one specific type of kernel matrix. The body chapters of the thesis (Chapters 2–5) otherwise use the precise term appropriate to each analysis.

Covariance, or more generally, kernel matrices can provide a simple and principled way to capture channel interaction structure in multi-channel recordings. By encoding second-order statistics, representations based on such matrices capture patterns of co-activity that can be sensitive to distributed brain states while being relatively insensitive to absolute signal scaling.

This makes covariance-based features particularly appealing for fNIRS, whose measurements often show large inter-subject or session-to-session variability.

Importantly, covariance matrices are not generic Euclidean feature vectors: they belong to the space of symmetric positive-definite (SPD) matrices, which has a non-Euclidean geometric structure. Standard Euclidean operations on covariance matrices, such as averaging or distance computation, can distort this structure and potentially degrade downstream analyses. This justifies the use of analytical frameworks that explicitly respect the geometry of covariance representations.

The appropriate mathematical framework for working with data that lie on curved manifolds, such as the space of SPD matrices, is Riemannian geometry. In Riemannian geometry, operations such as distances, means, and projections are defined directly on the manifold. This appropriately preserves the structure of representations. Over the past decade, Riemannian geometry-based methods have been successfully used in electroencephalography (EEG), where covariance matrices of sensor signals have been shown to serve as robust descriptors of brain states in brain-computer interface BCI and cognitive neuroscience applications (Barachant et al., 2012, 2013; Congedo et al., 2017).

EEG and fNIRS share essential structural similarities: both are multi-channel recording techniques in which covariance and co-activation across sensors provide a natural and informative representation of distributed neuronal activity. Despite this similarity in data representation, Riemannian geometry has not been utilized in fNIRS research to the same extent. This is a striking gap if one considers the growing emphasis on systems-level and distributed interpretations of hemodynamic signals and the demonstrated benefits of such geometry-aware methods in electrophysiology.

Riemannian Geometry for fNIRS

The first fNIRS project (Chapter 2) presented in this thesis addresses this gap by introducing Riemannian geometry-based machine learning models for fNIRS data analysis and classification. Based on previously established approaches from (BCI)-EEG research, the covariance matrices are constructed from the preprocessed fNIRS time series and subsequently manipulated on the SPD manifold. Given the unique structure of fNIRS signals, special attention is given to the representation of the two dependent variables measured with fNIRS: Oxygenated hemoglobin (HbO) and Deoxygenated hemoglobin (HbR). The first paper in this thesis introduces block-diagonal matrices that allow for independent hyperparameter tuning for HbO and HbR data separately (Näher et al., 2025). This approach underlines that HbO and HbR are separate, although correlated, signals that can both individually contribute to classification accuracy.

This methodological advancement shows that Riemannian classifiers can robustly decode internally generated brain states from fNIRS recordings and outperform classical Euclidean-feature approaches. Besides improvements in classification accuracy, the framework provides a conceptually coherent link between systems neuroscience and the mathematical structure of the data, emphasizing interactions rather than isolated activations.

Disorders of Consciousness: diagnostic challenges and the role of brain-based assessment

Disorders of consciousness (DoC) are clinically defined by disruptions in arousal and/or awareness after severe brain injury. DoC are commonly subdivided into three categories: coma, in which both arousal and awareness are absent; unresponsive wakefulness syndrome (UWS), characterized by preserved arousal without detectable awareness; and the minimally conscious state (MCS), in which arousal is intact and awareness is present but inconsistent (Edlow et al., 2021; Laureys et al., 2010; Thibaut et al., 2019). While coma can be diagnosed reliably based on behavior alone, the distinction between UWS and MCS remain a challenging problem in neurology. Gold-standard assessments like the Coma Recovery Scale-Revised (Giacino et al., 2004) and the Glasgow Coma Scale (Teasdale & Jennett, 1974) rely on overt motor behavior to infer awareness. This, however, is a fundamental problem in DoC, because motor impairments, involuntary movements, or irregular circadian rhythms can potentially mask remaining cognitive ability, which leads to misdiagnosis rates of up to 40% for MCS patients that are classified as UWS (Schnakers et al., 2009). The consequences of these misdiagnoses are substantial: MCS patients generally have a more favorable prognosis (Coleman et al., 2009), may retain the capacity for BCI-mediated communication (Schiff et al., 2024), and face different ethical considerations regarding continued treatment than patients without residual awareness (Andrews et al., 1996; Childs et al., 1993). These behavioral-based diagnostic limitations have motivated the development of brain-based diagnostic approaches. When covert mental imagery tasks were combined with non-invasive neuroimaging, some patients demonstrated deliberate modulation of brain activity without any motor output present (Monti et al., 2010; Owen et al., 2006). The most extensively studied modalities in this context are fMRI and EEG (Bardin et al., 2011; Claassen et al., 2019; Cruse et al., 2011; Galiotta et al., 2022). However, each of these modalities carries practical drawbacks as already outlined above. The portability, tolerance for movement, and simultaneous measurement of HbO and HbR make fNIRS particularly well suited to clinical settings where fMRI is impractical and EEG signal quality might be compromised. Despite this potential, existing fNIRS-based DoC paradigms often rely on canonical hemodynamic response functions, predefined imagery tasks with long encoding durations, and single-channel or channel-averaged analyses. These approaches might limit sensitivity to individual differences in brain-injury patterns and hemodynamic response profiles.

Clinical proof of principle

The second fNIRS project (Chapter 3) applies the previously introduced Riemannian geometry framework to the diagnostic gap in DoC outlined above. In a proof-of-principle study with healthy controls, participants simulated responsive and unresponsive states to mimic MCS and UWS. Covariance-based classification, as established in Chapter 2, was used to decide, at the individual participant level, whether the task instructions had been followed. A successful discrimination between command following and rest served as the marker of preserved awareness. The results demonstrate that covariance structure in fNIRS signals carries diagnostically relevant information, and that geometry-aware, systems-level analyses are not only theoretically motivated but also practically feasible in applied as well as clinical settings (Bastian et al., 2025).

fNIRS summary

Taken together, the fNIRS works in this thesis further establish fNIRS as a useful systems-neuroscience tool. Its full potential is best realized through analysis methods that use channel interaction-based representations. By transferring and adapting Riemannian approaches from EEG to fNIRS, the fNIRS studies in this thesis contribute both methodological advances, as well as empirical evidence that distributed hemodynamic interactions encode meaningful brain states with specific relevance for clinical applications. This perspective builds a foundational pillar of the broader thesis, which seeks to infer latent neuronal states from indirect or constrained measurements across different domains.

Part II: Saccade rhythmicity during active visual sampling

Vision is not a passive process. Rather than simply receiving incoming stimuli, primates actively reposition their eyes to selectively sample behaviorally relevant information from their visual environment. Saccadic eye movements shift the fovea to successive locations of interest, creating a sequence of high-acuity snapshots. During natural viewing conditions, these movements occur repetitively, with inter-saccadic intervals (ISIs) on the order of a few hundred milliseconds, resulting in saccade rates of approximately 2–6 Hz (Otero-Millan, Macknik, et al., 2013). This type of repetitive sampling bears a striking resemblance to other repetitive active-sensing behaviors observed across different mammals, such as whisking and sniffing in rodents. These active sampling behaviors have been shown to be highly rhythmic (Kleinfeld et al., 2006; Takatoh et al., 2022), raising the question whether saccade generation could be governed by a similar intrinsically rhythmic process. However, despite decades of research into saccade timing, the rhythmicity of oculomotor system output has recently been challenged (Amit et al., 2017).

Temporal structure of saccade sequences

The temporal structure of saccade sequences has been characterized along two complementary lines of evidence. The first line of evidence concerns the temporal regularities of immediately successive saccades in a sequence. Early work has shown that successive saccades cannot be initiated arbitrarily close in time: a post-saccadic refractory period imposes a hard floor on the

inter-saccadic interval, which leads to a minimum duration between successive saccades (Hallett & Adams, 1980; Hallett & Lightstone, 1976). This "dead zone" is not just a motor constraint but rather reflects the minimum time required to disengage fixation-related activity, reset the saccade generator, and subsequently allows a new motor plan to build toward a threshold. Combining this evidence with the characteristically skewed but unimodal shape of the ISI distribution, the temporal signature suggests that saccade generation exhibits distinct temporal structure beyond what would be expected from a memoryless and purely stochastic process.

Carpenter and colleagues proposed the LATER (Linear Approach to Threshold with Ergodic Rate) model, in which saccadic latency reflects the time for a linearly accumulating decision signal to reach a fixed threshold at a stochastically varying rate. This model naturally produces the right-skewed reaction time distributions that are empirically observed (Carpenter & Williams, 1995; Noorani & Carpenter, 2016). Although this model was originally formulated for stimulus-driven saccades, the same accumulation-to-threshold logic has been applied to the ISI distributions which have been observed during natural viewing conditions (Tatler et al., 2017). Bompas & Sumner (2011) extended this view to neuronal field models and showed that competitive fixation-saccade dynamics further shape the tail of the distribution. Taken together, these frameworks converge on a view that the ISI distribution is primarily a product of stochastic accumulation-to-threshold dynamics which are gated by a post-saccadic recovery period. Importantly, although these models were not intended to address rhythmicity as such, they implicitly contain a rhythmic component: each accumulation-reset cycle is functionally analogous to one period of an oscillator, with the period determined by the drift rate, noise, and effective distance to threshold.

Perhaps the clearest empirical evidence for an intrinsically rhythmic saccade generator has been made by Otero-Millan and colleagues across a series of studies. By pooling ISIs from both fixational microsaccades and larger exploratory saccades, Otero-Millan et al. (2008) showed that inter-event distributions retained a consistent shape across amplitude classes, with a Gaussian mode at around 150 ms and a long right tail which accommodates occasional extended intervals. The shared distribution shapes across saccades of different sizes suggest that the saccade timing is governed by the same underlying generation mechanism. Evidence from pathological conditions provides a further sign for an underlying rhythmic saccade generation. In square-wave jerks, which are involuntary conjugate saccade pairs that briefly move the eye away from and again towards the current fixation location, the interval between the two saccades clusters tightly around 200 ms (Otero-Millan et al., 2011; Otero-Millan, Schneider, et al., 2013). This can be interpreted as an oculomotor system that is effectively stripped of behavioral and cognitive demands, and therefore shows strong periodic structure in saccade generation, suggesting that a rhythmic kernel is built into the recovery dynamics of the saccade generator itself. Yet, its regular expression in healthy viewing might be masked by slow rate variability, top-down control signals, or other sources of ISI dispersion that accompany free exploration.

The confound of rate non-stationarity

Despite the converging evidence for local temporal structure reviewed above, the question of whether saccade generation during free viewing reflects a genuinely rhythmic process has received surprisingly little attention. To my knowledge, the question has been addressed head-on in only a single study. Amit et al. (2017) compared two candidate explanations of free-viewing saccade generation: a self-paced process governed only by first-order dynamics such as the post-saccadic refractory period, and a strict fixed-frequency oscillator acting as a central pattern generator. Using spectral measures, autocorrelation analyses, and ISI shuffling tests, Amit et al. (2017) concluded that a strict, continuously phase-advancing pacemaker was not needed to explain the characteristic spectral theta peak of saccade sequences. Instead, the spectral peak could be parsimoniously explained by the refractory period alone. However, their analyses did not rule out a more flexible locally rhythmic generator, and the question of whether saccade timing during natural viewing reflects an intrinsically rhythmic process therefore remains open.

A key source of this ambiguity lies in a property of natural viewing saccades that is easily overlooked in the design of their rhythmicity analyses: saccade rate is not stationary over time. Rather, it fluctuates across trials, drifts gradually within prolonged viewing epochs, and varies systematically with stimulus content and task demands (Otero-Millan et al., 2008). Under these conditions, the empirically measured ISI distribution reflects a marginal distribution that is aggregated across many local rate regimes, with each potentially generating intervals of different typical durations. A fundamental and under-appreciated consequence of this aggregation is that mixing regular waiting times across variable rates produces heavy-tailed distributions, regardless of how regular the underlying local process within one rate regime is. Because heavy tails have previously been read as a signature of stochastic, non-rhythmic generation (Amit et al., 2017), their ubiquity in saccade timing data has often been taken as direct evidence against underlying rhythmicity. But crucially, this inference collapses two distinct sources of temporal variability into one: local scattering in the timing of individual saccades, and the slow modulation of saccade rate itself.

This confound extends beyond the shape of the ISI distribution to essentially every diagnostic that has been used to assess saccade rhythmicity. Spectral peaks computed over the marginal distribution are systematically attenuated by averaging across rate regimes, because the oscillatory signature at any given local rate is smeared across frequencies when the rate drifts. ISI shuffling, which was intended to test whether temporal dependencies extend beyond immediately adjacent saccade intervals, carries the implicit assumption of stationarity and can fail to detect the higher-order structure that rate non-stationarity introduces into the ISI sequence. This leads to the conclusion that the standard analytical toolkit for saccade rhythmicity may be systematically biased toward negative conclusions about underlying rhythmicity: a rhythmic generator might not be absent but just occluded by a source of variability in saccade timing that operates on a slower timescale than the one these tools are designed to probe.

Analytical contribution: decomposing saccade timing

The first saccade study presented in this thesis addresses this confound by developing an analytical approach that separates local saccadic rhythmicity from slow saccade rate fluctuations. Based on tools from neuronal spike train analysis (Holt et al., 1996), saccade timing can be modeled as a doubly stochastic gamma-renewal process. In such a model, a "locally regular" generator is being modulated by a slowly varying rate. This model is then used for the analysis of human, macaque, and marmoset saccade timing data in matched visual foraging tasks (Näher et al., 2023). Two previous observations that have historically undermined the case for saccadic rhythmicity (the right-skewed shape of the ISI distribution and the absence of sharp spectral features) turn out to be compatible with an essentially periodic local generator once the rate is allowed to change over time. To measure rhythmicity without being affected by this rate-driven variability, the study proposes the rate-corrected rhythmicity index (RCRI). This index rescales successive intervals in a way that is largely immune to slow drifts in saccade rate. When the RCRI is applied to data from the three primate species, a consistent pattern becomes evident: the faster the oculomotor system samples the environment, the more regular the local timing between saccades becomes. This observation is highly compatible with the formulation of a drift-to-threshold model (Carpenter & Williams, 1995; Noorani & Carpenter, 2016) where a faster rise-to-threshold compresses the noise that the process accumulates between successive saccades.

Causal test: perturbing the cortical oculomotor network

The descriptive decomposition outlined in the first saccade study (Chapter 4) gives rise to a concrete mechanistic prediction. In a drift-to-threshold model of saccade generation, the mean saccade is determined by the average distance it takes for the generative process to reach its threshold. However, within a given rate regime, local regularity/rhythmicity of the process reflects the amount of noise in this rise-to-threshold between successive saccades. The source of noise to this process could be within the generator itself or come from different and potentially cortical areas that provide top-down control over visual sampling. The frontal eye field (FEF) is a particularly compelling candidate. As a central node of the dorsal attention network (DAN), it is continuously engaged in selecting the next saccade target, updating priority representations of the visual scene, and driving movement-selective neurons toward threshold (Hanes & Schall, 1996; Schall, 2002). Ongoing cortical planning of this kind is expected to inject variability into the downstream saccade circuits, adding noise to the local generator on top of its intrinsic recovery dynamics.

This line of reasoning yields a testable prediction. If cortical deliberation is indeed a source of timing noise, then inhibition of the relevant cortical areas should allow the intrinsic recovery dynamics of the generator to dominate again. This would lead to the local timing structure being largely shaped by the recovery dynamics and therefore increasing rhythmicity even when mean rate is unchanged. If such a dissociation between rate and rhythmicity can be observed

empirically, it provides compelling evidence that the observed, marginal saccade rhythmicity is a compound of a baseline rhythmicity that is likely generated by intrinsic recovery dynamics and additional injected noise from higher cognitive areas. The second saccade study in this thesis (Chapter 5) evaluates this prediction causally, using continuous theta-burst stimulation (cTBS), an offline transcranial magnetic stimulation protocol that transiently reduces cortical excitability and creates a "virtual lesion" (Huang et al., 2005; Nyffeler et al., 2006). cTBS is applied to right FEF while participants perform a visual search paradigm. Between search blocks, resting state EEG is acquired at multiple time points before and after the cTBS intervention. Source-localized spectral and connectivity analyses are used to characterize stimulation-induced changes in cortical excitability, oscillatory dynamics, and in the functional coupling between right FEF and the broader DAN. By combining the rate-aware analytical framework developed in the first study with a causal perturbation targeting a candidate cortical source of timing variability, this study seeks to fill the gap between descriptive timing statistics and the circuit-level control variables that shape them.

Saccade rhythmicity summary

Taken together, the saccade work in this thesis establishes that primate saccade generation during natural viewing contains a genuine local rhythmic component. This component is not visible in conventional analyses due to slow rate fluctuations masking it. However, it becomes clearly detectable once rate fluctuations are accounted for. The first study develops the analytical decomposition and shows, across the three most used primate species in neuroscience, that local rhythmicity scales monotonically with saccade rate, consistent with a drift-to-threshold generator. The second study further tests this framework with a causal manipulation by transiently inhibiting right FEF with cTBS, linking cortical deliberation to the cycle-to-cycle variability that shapes local timing. Together, these studies position saccade generation alongside other rhythmic active-sensing behaviors.



2

Chapter 2

Riemannian Geometry Boosts Functional Near-Infrared Spectroscopy-Based Brain-State Classification Accuracy

This chapter is based on the following publication:

Näher, T., Bastian, L., Vorreuther, A., Fries, P., Goebel, R., & Sorger, B. (2025). Riemannian geometry boosts functional near-infrared spectroscopy-based brain-state classification accuracy. *Neurophotonics, 12*(4), 045002.

DOI: https://doi.org/10.1117/1.NPh.12.4.045002

Abstract

Background: Functional near-infrared spectroscopy (fNIRS) has recently gained momentum as a reliable and accurate tool for assessing brain states based on the vascular response to neural activity. This increase in popularity is due to its robustness to movement, non-invasive nature, portability, and user-friendly application. However, compared with other hemodynamic functional brain-imaging methods such as functional magnetic resonance imaging (fMRI), fNIRS is constrained by its limited spatial resolution and coverage with a particularly limited penetration depth. In addition, due to comparatively fewer methodological advancements, the performance of fNIRS-based brain-state classification still lags behind more prevalent methods such as fMRI.

Methods: We introduce a classification approach grounded in Riemannian geometry for the classification of kernel matrices, leveraging the temporal and spatial relationships between channels and the inherent duality of fNIRS signals, specifically oxygenated and deoxygenated hemoglobin. For the Riemannian-geometry-based models, we compared different kernel matrix estimators and two classifiers: Riemannian Support Vector Classifier and Tangent Space Logistic Regression. These were benchmarked against four models employing traditional feature extraction methods. Our approach was tested on seven participants in two brain-state classification scenarios based on the same fNIRS dataset: an eight-choice classification, which includes seven established plus an individually selected imagery task, and a two-choice classification of all possible 28 two-task combinations.

Results: This approach achieved a mean eight-choice classification accuracy of 65%, significantly surpassing the mean accuracy of 42% obtained with traditional methods. In addition, the best-performing model achieved an average accuracy of 96% for two-choice classification across all task combinations, compared with 78% with traditional models.

Conclusion: To our knowledge, we are the first to demonstrate that the proposed Riemannian-geometry-based classification approach is both powerful and viable for fNIRS data, substantially increasing the accuracy in binary and multi-class classification of brain activation patterns.

Introduction

Functional near-infrared spectroscopy (fNIRS) is an emerging functional neuroimaging technique that has recently gained momentum due to its non-invasive nature, portability, silent data acquisition, robustness against movement artifacts, and favorable trade-off between spatial and temporal resolution (Klein et al., 2024; Kohl et al., 2020; Soekadar et al., 2021). FNIRS detects hemodynamic responses, which are also measured by functional magnetic resonance imaging (fMRI) and have been linked to changes in neuronal activity (Logothetis et al., 2001). Both fMRI and fNIRS assess neuronal activity indirectly via the blood-oxygen level-dependent (BOLD) signal. The BOLD signal is especially linked to neuronal gamma oscillations (Koch et al., 2009; Scheeringa & Fries, 2019). Intriguingly, the BOLD signal reflects changes in both gamma-band activity and neuronal firing rates when they are driven by external stimuli; yet when changes are due to internal state fluctuations, they are captured by BOLD specifically for gamma and much less for firing rates (Niessing et al., 2005). This makes gamma oscillations a crucial marker for studying brain activity, particularly in the context of internal state changes (Koch et al., 2009; Scheeringa & Fries, 2019).

Unlike fMRI, fNIRS measures two distinct hemoglobin signals: oxygenated (HbO) and deoxygenated (HbR) hemoglobin. Although these signals are naturally anticorrelated (Obrig & Villringer, 2003), they are not merely redundant; rather, they can provide complementary information about brain activation (Tam & Zouridakis, 2015). Despite measuring two separate signals simultaneously, the accuracy of fNIRS-based classification of brain states has yet to reach the level of fMRI. This disparity could partially be attributed to the fact that fMRI benefits from more advanced and well-developed data-analysis techniques due to its more frequent use compared with fNIRS. Consequently, the full potential of fNIRS data analysis has not yet been realized, given its historically less developed and less widely known status. Although fMRI can classify brain states by utilizing spatial differentiation of multivariate activation patterns across the whole brain (Gardumi et al., 2016), fNIRS is constrained by its limited spatial resolution and coverage, particularly its limited penetration depth. As a result, not all classification methods suited for fMRI data may perform as effectively when applied to fNIRS data.

The current methods for fNIRS-based brain-state classification are often adapted from other functional neuroimaging techniques but are not specifically tailored to the unique properties of fNIRS signals. Traditional classifier approaches include linear discriminant analysis (LDA), support vector classifier (SVC), random forest (RF), and logistic regression (LR), which typically rely on channel-based features such as mean activation, channel variance, regression slope, and zero crossings (Herff et al., 2013; Naseer & Hong, 2015; Noori et al., 2017). Although more advanced methods, such as convolutional and long short-term memory (LSTM) neural networks, have been employed to improve classification accuracy, such deep-learning approaches require large amounts of training data, which are often not available (Eastmond et al., 2022).

Given these limitations, there is a need to develop methods specifically tailored to the unique properties of the fNIRS response. In practice, this involves emphasizing the dual nature of fNIRS signals by utilizing the complementary information potentially provided in HbO and HbR data. Advancing data analysis by considering the unique fNIRS-signal properties could bridge the gap between the practicality of fNIRS and the precision seen in fMRI, making fNIRS a more viable alternative for brain-state classification. The aim of this study was to introduce the application of Riemannian-geometry-based models to fNIRS data. To achieve this, we combined HbO and HbR signals in block diagonal channel kernel matrices and subsequently applied Riemannian-geometry-based classification models. Instead of relying on features from individual channels, Riemannian-geometry-based models additionally utilize relationships between channels. To our knowledge, we are the first to demonstrate that the classification of fNIRS data using Riemannian geometry offers a promising alternative to other machine-learning methods.

Riemannian geometry machine-learning models work by differentiating brain states based on different spatial co-activations. The framework offers well-established tools that allow for the independent manipulation and therefore tuning of separate signals within the same model. This allows us to effectively capture spatial co-activation patterns separately for both HbO and HbR with only one model. To further optimize task selection for classification applications, we also recorded participant ratings regarding task ease and comfort, as tasks perceived more favorably might inherently provide clearer neural activation patterns. Our approach outperforms traditional methods in both eight- and two-choice classification scenarios. Although Riemannian-geometry-based machine learning is a well-established technique in other classification fields (Barachant et al., 2010; Congedo et al., 2017; Formont et al., 2013; Harandi et al., 2018; Pennec et al., 2006; Tuzel et al., 2006; Zhai et al., 2021), we show that it is also a powerful and viable option for fNIRS data.

Methods

Participants

Seven healthy participants (mean age: 32.1 +/- 10.6 years, three females) were included in this experiment (Supplementary Table 1). All participants were either externally recruited or students and staff members of the Faculty of Psychology and Neuroscience at Maastricht University. Written informed consent was acquired before the experiment. The local ethics committee approved the experiment (Ethics Review Committee Psychology and Neuroscience). Upon successful completion of the experiment, participants received monetary compensation.

,

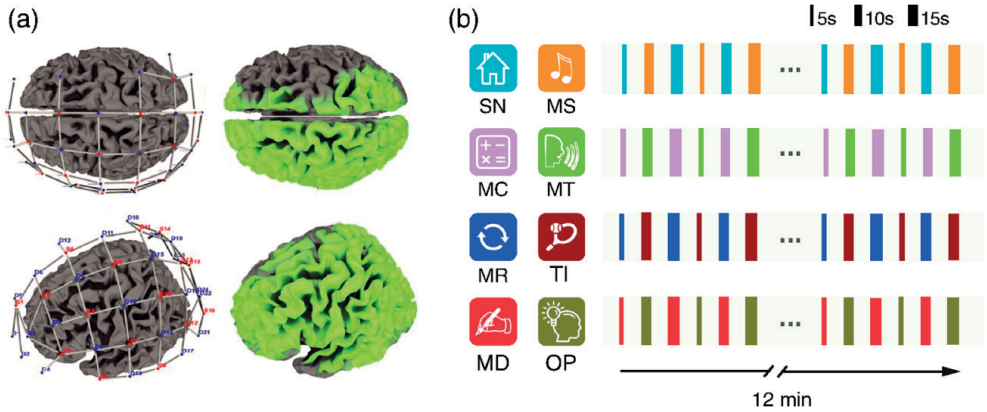

Figure 1. Optode array with resulting NIR light-intensity map and experimental procedure.
(a) Optode array (left) and its resulting near-infrared (NIR) light-intensity map (right) that indicates the brain coverage (light green regions). **(b)** Example structure of task pairs used in the four functional runs for one participant. The task pairings varied across participants. Each run started with a 1 min rest period. Task pairs were presented in alternation for 24 trials with 20s rest periods after each task trial. Encoding durations varied between 5 s, 10 s, and 15 s. This pattern was repeated four times in each run and resulted in a total duration of 12min per run. Abbreviations: MC: Mental Calculation, MD: Mental Drawing, MR: Mental Rotation, MS: Mental Singing, MT: Mental Talking, OP: Own Paradigm, SN: Spatial-Navigation Imagery, TI: Tennis Imagery.

General procedure

Before the experiment, participants received instructions (see Table 1) via email, including short but precise mental-task descriptions to train task performance beforehand. Participants were asked to familiarize themselves with this information and devise their "own (mental) paradigm" (OP): A suitable mental task or process that was not included in the provided task set. The OP task should be familiar to them, and they should feel comfortable executing it upon being cued. We included the OP task to probe whether a participant-defined task would be more effective compared to predefined tasks as it may enhance motivation and task engagement, and thereby imagery success. During the fNIRS session, each participant was seated in a darkened and sound-attenuated room. The fNIRS cap was placed on the participant's head, and all optodes were adjusted to ensure orthogonality to the scalp and stable optode-scalp contact. During optode placement, participants were presented with the same stimulation (task cues) used for the experiment to practice mental task performance until they felt comfortable. Participants were also instructed to get into a comfortable position and to relax their jaw and facial muscles during the recording session. The fNIRS session lasted for approx. 1.5 hours including participant preparation. At the end of the session, participants rated experienced pleasantness and easiness levels of all eight mental task performances on a 10-point Likert scale (i.e., 1 = extremely unpleasant/extremely difficult to 10 = extremely pleasant/extremely easy). Ratings were not included after each run to ensure that participants performed comparative evaluations across runs and to maintain a smooth and uninterrupted task flow.

Experimental runs

Participants performed four runs with a total of eight mental tasks (Fig 1b). Mental tasks included Mental Talking (MT), Spatial-Navigation Imagery (SN), Mental Drawing (MD), Mental Singing (MS), Mental Calculation (MC), Mental Rotation (MR), Tennis Imagery (TI), and the previously devised Own Paradigm (OP). All tasks except for OP have been shown to elicit differential brain activation patterns using fMRI (unpublished data from the lab).

Table 1. Mental Task Instructions.

Mental Task (Abbreviation)	Visual Stimulus	Instructions
Mental Calculation (MC)	cal	During mental calculation, please calculate multiplication tables ('the basics'). Perform calculating the multiple of 7, 8, or 9 up to the decuple (e.g., 7-14-21-28-35-42-49-56-63-70, 8-16-24…). Do that at a comfortable but consistent speed. Try to visually imagine the numbers. This might support your imagination. Start from the beginning if necessary.
Mental Drawing (MD)	dra	During mental drawing, please imagine drawing simple geometric figures (such as circles, triangles, cubes, etc.) or small contour drawings (e.g., a butterfly, star, car, tree, boat, or house). Do that with your dominant hand and at a comfortable but consistent speed. Try to imagine using a pen. This might support your imagination. Start from the beginning if necessary.
Mental Rotation (MR)	rot	During mental rotation, please imagine either a clock with rotating clock hands or a diver jumping from a tower into the water while he spins around several times. Only imagine the moment he is moving in the air. Thus, in your imagination, he should continuously spin around. Again, try to perform this mental task with a comfortable but consistent speed as long as it is indicated.
Mental Singing (MS)	sin	During mental singing, please covertly sing any song you know by heart, e.g., a nursery rhyme. Please sing in a comfortable but consistent speed. Try to hear yourself singing. This might support your imagination. Start with the song from the beginning if necessary.
Mental Talking (MT)	tal	During mental talking, please covertly recite a text you know by heart (e.g., the prayer 'Our Father' or any poem you know). Try to hear yourself talking. This might support your imagination. You might have to start from the beginning in case the text ends before the end of the mental task. Start from the beginning if necessary.

Spatial-Navigation Imagery (SN)	nav	During spatial-navigation imagery, you should imagine to 'go' through your house/apartment and look into the different rooms for a moment (e.g., 2 s). Do that at a comfortable but consistent speed. Try to really imagine vividly the various three-dimensional (3D) scenes. The order of the rooms does not matter. However, try to continuously perform this mental task (thus continuously try to imagine the 3D scene of the particular room). After having looked into all rooms of your house/apartment, immediately start from the beginning and perform the spatial navigation task as long as it is indicated.
Tennis Imagery (TI)	ten	Imagine playing tennis against a wall (forehand – backhand – forehand – etc.). Focus on the motor aspect (do not visualize the scene). Keep the playing speed constant and perform the tennis imagery as long as it is indicated.

Note. The table presents an overview of the mental tasks performed by participants in the current fNIRS study, including abbreviations in parentheses. The visual stimulus presented during the experimental session was limited to three characters (second column). Detailed instructions on how to perform each task were provided to participants prior to the experimental session.

Participants were instructed to engage in the mental task cued on a computer screen for the duration that the cue was shown. To instruct participants when to engage in mental imagery, three letters of the mental task name were displayed (e.g., "dra" instead of "mental drawing"; see other abbreviations used in Table 1). This was meant to keep visual stimulation similar across tasks and avoid visual stimulation-based differences in cerebral activation. Instructions were presented for 2 s, followed by a fixation cross. Participants were told to engage in the mental task displayed on the screen with the greatest possible consistency. Each trial was followed by a resting period, which was instructed by displaying "res", and during which participants should not engage in a specific thought or task. Each participant performed four mental imagery runs. For each run, two of the eight mental tasks were randomly selected, without replacement, such that across the four runs, each of the eight tasks were selected once. Within each run, the two selected tasks were performed in alternating trials (Fig 1b). Each trial had a duration that was randomly chosen to be 5 s, 10 s or 15 s, while the duration of the following rest period was always 20 s. Each task was performed for a total of twelve trials per run, resulting in 24 trials per run and 96 trials per participant (Fig 1b). Across participants, the pairings of tasks were pseudo-randomized to avoid the same pairings across participants.

Data acquisition

Data were acquired with a NIRScout 16x24 system (NIRx Medizintechnik GmbH, Berlin, Germany) (Fig 1a). The twenty-four detectors and 16 sources were positioned according to the international 10-20 Electroencephalography (EEG) system. The following list contains the optode positions according to the 10-20 system. Source positions were over AFz, F5, F1, FT7, FC3, FCz, C5, C1, TP7, CP3, CPz, P5, P1, P2, POz and O1. Detectors were positioned over FPz, FP1, AF3, AF7, AF4, Fz, F3, FC1, FC5, C3, Cz, FC2, T7, CP5, CP1, CP2, P7, P3, Pz, P4, PO7, PO3, PO4, and Oz. Optodes provided full coverage of the left hemisphere and targeted coverage of task-relevant regions in the right hemisphere. This montage was chosen to include areas where we expected to observe mental task-related brain activity. Because most of our mental imagery tasks elicit predominantly left-hemispheric activation including language comprehension and production areas, we sampled large parts of the left hemisphere. We also included bilateral parietal coverage to capture activations associated with mental rotation and spatial navigation, which reliably engage parietal cortices in both hemispheres (Cohen et al., 1996; Epstein et al., 2017). Data were acquired using NIRStar 15.2 (NIRx Medizintechnik GmbH, Berlin, Germany) at 3.47 Hz.

Data preprocessing

All preprocessing steps were executed using *Satori* (v2.0). According to the optode array (Fig 1a), raw data were preprocessed in 62 channels for oxyhemoglobin (HbO) and 62 channels for deoxyhemoglobin (HbR). Data were trimmed to begin 5 s before the first mental-task trial onset and end 15 s after the last trial offset. The raw time course values were converted to HbO and HbR concentration values using the modified Beer-Lambert law. Subsequent motion correction included spike removal (ten iterations, 5 s lag, 3.5 threshold, 0.5 influence) with monotonic interpolation and Temporal Derivative Distribution Repair (TDDR), preserving high-frequency components (Fishburn et al., 2019). Short-separation channels are source–detector pairs placed close together (approximately 1 cm apart), such that the detected light passes only through the scalp and skull and does not reach the brain. Their signal therefore reflects systemic physiological noise, such as heartbeat, respiration and skin blood flow, which can then be regressed out of the regular channels (short-separation regression). Due to hardware limitations, we did not obtain time courses from short channels, which is why we were not able to perform short-separation regression to correct for systemic noise such as heartbeat, respiration and skin-blood flow (Haeussinger et al., 2014; Keshmiri et al., 2019; Kohno et al., 2007; Sato et al., 2016; Tak & Ye, 2014; Q. Zhang et al., 2009). Instead, a global component regression based on principal component analysis (PCA) was performed, removing the first component. This approach assumes that systemic physiological noise affects all channels in a similar way and is therefore captured by the first principal component across channels, which we removed from the data (Virtanen et al., 2009). Finally, the signals were filtered using Gaussian low-pass (0.4 Hz) and Butterworth high-pass (0.01 Hz) and z-scored. For each individual task trial, we cut the data from cue onset to 10 s post cue. Note that during the recording of the task, participants performed

mental imagery for either 5 s, 10 s, or 15 s, followed by 20 s of rest. To keep trials at the same length, we only utilized the first 10 s post cue for classification analysis. We estimated several features per channel, including mean, variance, peak, zero-crossings, skew, and kurtosis. These features were then input into the traditional classifiers. We chose to keep the length of windows the same, regardless of task performance duration to avoid introducing a bias in covariance estimation due to different amounts of data. This ensured that covariance and kernel estimates were based on comparable amounts of data across all trials, independent of whether the underlying task duration was 5, 10, or 15 s. Hyperparameters related to the classifiers underwent hyperparameter tuning to optimize their performance. For the Riemannian methods, different kernel matrices were estimated from the 62 HbO and HbR channels. The estimation of these kernel methods was according to several different covariance and (non)-linear kernel functions which were also considered hyperparameters of the Riemannian models. The resulting matrices served as features for the Riemannian classifiers, which also underwent hyperparameter tuning.

Classifiers

We compared 14 classifiers: ten Riemann-based models and four traditional models. The traditional models were Linear Discriminant Analysis (LDA), Logistic Regression (LR), Support Vector Classifier (SVC), and Random Forest (RF). We computed standard features (mean, variance, peak, zero-crossings, skew, and kurtosis) for these models as described in the fNIRS literature (Naseer & Hong, 2015). The hyperparameters for each model can be found in the Supplementary Table 2. Hyperparameters were optimized using an exhaustive grid search, in which all possible parameter combinations were evaluated using 5-fold cross-validation. Hyperparameter tuning was performed separately for each participant on that participant's own data, and because tuning and evaluation were not nested, the reported accuracies may be slightly optimistic, although the same procedure was applied to all models.

Kernel matrix estimation

Terminology. In this chapter, we use "kernel matrix" as the umbrella term for any channel-by-channel matrix in which each entry quantifies how similar the time series of two channels are. The function that computes this similarity is called the kernel function. A covariance matrix is one specific type of kernel matrix, in which similarity is measured as the linear co-variation of two mean-centred channel time series. It can be computed with different estimators, such as the standard sample estimator, shrinkage estimators, or robust M-estimators. The correlation matrix is its normalized form. In contrast, non-linear kernel matrices, such as radial basis function (RBF), Laplacian, polynomial, or cosine kernels, measure similarity in a way that also captures non-linear relationships between channel time courses. All of these matrices are symmetric and, after shrinkage regularization, positive definite, which means they can be analyzed with the same Riemannian tools. Accordingly, in the "block covariance" models both blocks are covariance or correlation matrices, in the "block kernel" models both blocks are non-linear kernel matrices,

and in the "super kernel" models each block can be either type, chosen independently for HbO and HbR.

For EEG-based BCIs, classification based on covariance matrices is standard practice (Barachant et al., 2010, 2013; Yger et al., 2017). These matrices, specifically termed "channel covariance matrices," are derived from the covariances between channel pairs. Covariance measures the extent to which two time series vary together: a positive covariance indicates that increases in one time series are matched by increases in the other, whereas a negative covariance suggests that an increase in one corresponds to a decrease in the other. However, note that a covariance of zero does not necessarily denote independence between the time series; rather, it indicates that there is no linear relationship between them.

We used several different kernel functions to obtain the kernel matrices between channels, employing the following estimators: the correlation coefficient matrix (“corr”), the numpy-based covariance matrix (“cov”), Huber’s M-estimator based covariance matrix (“hub”), the Ledoit-Wolf shrunk covariance matrix (“lwf”), the oracle approximating shrunk covariance matrix (“oas”), the Schaefer-Strimmer shrunk covariance matrix (“sch”), the scikit-learn empirical sample covariance matrix (“scm”), Tyler’s M-estimator based covariance matrix (“tyl”), the polynomial kernel (“poly”), which adds non-linearity by raising the dot product to a specified power, the radial basis function (RBF) kernel (“rbf”), which measures similarity based on the distance between input vectors; the Laplacian kernel (“laplacian”), which is similar to the RBF kernel but uses the L1 norm for distance measurement, and the cosine kernel (“cosine”), which measures the cosine of the angle between input vectors. All kernel functions were implemented using scikit-learn and pyriemann using default parameters.

Let K_{HbO} be the kernel matrix for the HbO channels, and K_{HbR} for the HbR channels, where the kernel function $\kappa(\cdot,\ \cdot)$ computes the relationship between pairs of channels. Formally, for an input time series $X \in R^{c \times t}$, composed of c channels and t time samples, the elements of the kernel matrices are defined as:

$$K_{i,j}^{\mathrm{HbO}} = \kappa\left(X[\mathrm{HbO}]_i, X[\mathrm{HbO}]_j\right)$$

$$K_{i,j}^{\mathrm{HbR}} = \kappa\left(X[\mathrm{HbR}]_i, X[\mathrm{HbR}]_j\right)$$

Here $X[\mathrm{HbO}]_i$, and $\mathcal{X}[\mathrm{HbO}]_j$ denote the i-th channel time series for HbO and HbR, respectively. The matrices K_{HbO} and K_{HbR} capture the intrinsic channel-to-channel relationships within their chromophore.

Instead of classifying HbO and HbR kernel matrices separately with different classifiers, we can combine them in a block diagonal matrix. A block diagonal matrix is a square matrix composed of smaller square matrices, or "blocks," along its diagonal and zero entries elsewhere.

The combined kernel matrix, K_{block}, structured as a block diagonal of K_{Hbo} and K_{HbR}, is defined as:

$$K_{block} = K_{Hbo} \oplus K_{HbR} = \begin{bmatrix} K_{Hbo} & 0 \\ 0 & K_{HbR} \end{bmatrix}$$

It is important to note that during hyperparameter tuning, we allow $\kappa(\cdot, \cdot)$ to be (non-linear) kernel functions as well as covariances with the different estimators, as described above. $\kappa(\cdot, \cdot)$ can also be different for HbO and HbR. To ensure the numerical stability of the kernel matrices, we additionally regularized of K_{HbO} and K_{HbR} separately before combining them to K_{block}.

This regularization was done via shrinkage of the matrices, allowing different shrinkage parameters for each chromophore. This approach emphasizes flexibility for each oxygenation type and allows the classifier to be tuned separately to HbO and HbR data while ensuring the SPD property of the matrices. We termed this kernel matrix estimation “super kernel”. A python implementation of our super kernel estimator can be found in the pyriemann documentation (called HybridBlocks: https://pyriemann.readthedocs.io/).

In contrast, a full kernel matrix estimation does not permit such tailored optimization, which can limit performance. For full kernel matrix estimation, the approach is essentially the same as described earlier, however, only one kernel function is picked for both HbO and HbR. These full kernel matrices also include the interactions between HbO and HbR channels, while these interactions are set to zero for the block diagonal matrices. This means that the full kernel matrices can contain more information than the block diagonal matrices. However, they also lack the ability to tailor kernels separately for HbO and HbR and occupy a higher dimensional manifold.

Co-activation comparison and mean kernel matrices

To test whether HbO and HbR show different co-activation patterns, we compared distances between the mean kernel matrices. For this, we first computed the mean kernel matrix for each task and for each participant, separately for HbO and HbR. Note that all mean kernel matrices were estimated with the Riemannian metric (Fletcher et al., 2004; Moakher, 2005). Next, we computed the pairwise Riemannian distances (Chevallier et al., 2021) between all pairs of tasks (n = 28). The resulting distance vectors for HbO and HbR were transformed into discrete probability distributions. Subsequently, we computed the Wasserstein distance between these vectors. We used the Wasserstein distance because it provides a natural metric for comparing distributions based on shape and magnitude, allowing an intuitive interpretation of differences

between HbO and HbR kernel matrices. The Wasserstein distance, also known as Earth mover's distance, can be thought of as the minimum "cost" required to transform one distribution into the other. The "cost" here is calculated based on the amount of probability mass that needs to be moved, and the distance it needs to be moved. Intuitively it represents a difference in shape between the two distributions. The logic is as follows: If HbO and HbR kernel matrices both contain the same information, the distribution of pairwise distances between the individual task pairs is similar. This would lead to a small Wasserstein distance. Likewise, if the two chromophores contain different information, and can hence differently differentiate between task pairs, the distance will be larger. We calculated the Wasserstein distance between HbO and HbR pairwise task distances for all participants and tested whether this distribution differed significantly from zero with a one-sample *t*-test.

Classification approach

The dataset used in this study consisted of eight different mental tasks with twelve trials each for a total of 96 trials per participant. Given the relatively small sample size, we opted against splitting data into separate training, testing, and validation sets. Instead, we employed a repeated and stratified k-fold cross-validation approach, which is a common approach in machine learning scenarios with small datasets. Specifically, we employed repeated (n = 5) 5-fold stratified cross-validation to compute the accuracy score for model comparisons. While this method is widely accepted and helps mitigate the variability in small datasets, it is important to note that it can sometimes overestimate the accuracy compared to using a separate test set (Hastie et al., 2009). However, we performed the same approach for both traditional models as well as the Riemann models. All computations were performed on a high-performance computing (HPC) cluster using SLURM, ACME, scikit-learn, and pyriemann.

Relative improvement score

We calculated a relative improvement score for Riemann models compared to traditional models for each participant. This score normalizes the improvement by considering the proximity of the baseline accuracy to the maximum possible accuracy (i.e., 1). The relative improvement score is defined as follows:

Let A_{baseline} be the accuracy of the baseline model, and A_{new} be the accuracy of the new model. The relative improvement score RIS is calculated using the formula:

$$RIS = \frac{A_{\text{new}} - A_{\text{baseline}}}{1 - A_{\text{baseline}}}$$

This formula scales the observed improvement by the remaining potential for improvement, given the baseline performance. The numerator, $A_{\text{new}} - A_{\text{baseline}}$, represents the absolute improvement in accuracy achieved by the new model. The denominator, $1 - A_{\text{baseline}}$ represents the maximum possible improvement from the baseline model's accuracy to perfect accuracy. By

using this relative improvement score, we account for the fact that achieving improvements is more challenging as the baseline accuracy approaches 1.

Statistical tests for model comparison

We compared the classification accuracies of traditional methods to the Riemann methods using paired *t*-tests. Since we used ten Riemann models and four traditional models, we calculated the average accuracy for each participant and each model type using a bootstrap procedure (n = 1000). The bootstrapping was applied to the cross-validation accuracies to obtain a more robust estimate of the distribution and to compute 95% confidence intervals. This allowed us to compute the mean of the bootstrapped accuracy distribution and then perform the t-tests on these bootstrapped means.

For the best-performing model for the 2-choice and 8-choice comparison, we determined whether the repeated stratified 5-fold cross-validation accuracy was above chance level. The theoretical chance levels for the 2- and 8-choice comparisons are 50% and 12.5%, respectively. However, these numbers are only valid for an infinite number of samples. To better estimate classification performance, we used a permutation-based approach where we shuffled the labels (n = 500) and retrained models on the shuffled labels (Combrisson & Jerbi, 2015). All analyses used a significance threshold of $P < 0.05$ unless stated otherwise.

Results

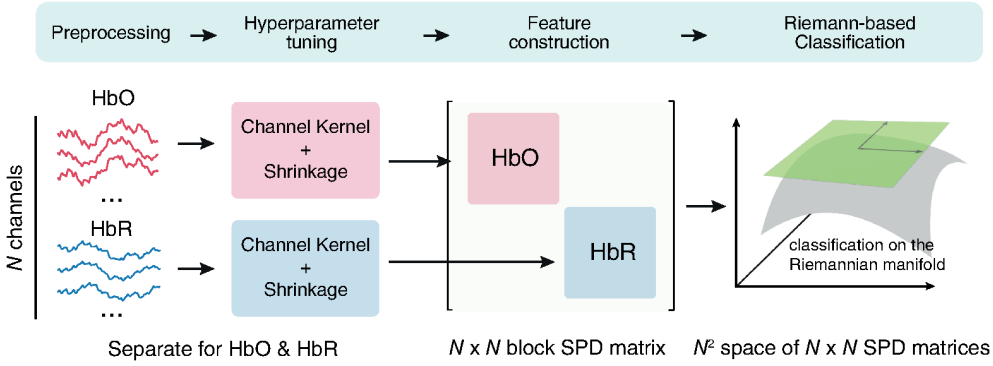

Figure 2. Processing pipeline of Riemannian fNIRS classifiers using block diagonal matrices as features. The data are first preprocessed according to standard pipelines. The subsequent hyperparameter tuning finds the best kernel functions and regularization for HbO and HbR separately. Each combination of hyperparameters is evaluated using a 5-fold stratified and repeated cross-validation (five repetitions). The resulting individual kernel matrices are then combined in a block diagonal matrix which serves as a feature for each trial. Several Riemannian-based classifiers were evaluated for classifying fNIRS data across participants.

Leveraging the inherent duality of fNIRS data with block diagonal matrices

To optimally use the information provided by both HbO and HbR, we designed models that use block diagonal matrices for classification. This has the advantage that the manifold on which the data are classified contains information from both HbO and HbR signals that the model can leverage. The pipeline for fitting and estimating these models is illustrated in Fig 2.

The process begins with standard fNIRS preprocessing steps (detailed in the Methods), during which the data are separated into HbO and HbR signals. After preprocessing, the HbO and HbR data are processed separately (Fig 2). For each chromophore, we compute a kernel matrix with different estimators according to the hyperparameter set. This separation allows the hyperparameter search to identify the optimal function for each chromophore, optimizing the sensitivity and overall performance of the model. Following the estimation, the HbO and HbR matrices are individually regularized using shrinkage methods. This regularization step is crucial for stabilizing the matrices, especially given the high dimensionality of the data. Once regularized, the two matrices are combined into a block diagonal matrix, which is then used as input for a classifier (Fig 2). This method offers a significant advantage by allowing the most efficient kernel and shrinkage parameters to be determined for each signal type individually.

We benchmarked 14e different classifiers in two different settings: Classifying eight tasks and classifying two tasks. For each setting, we evaluated four traditional machine learning models and ten Riemannian Geometry-based models. The traditional models were LDA, LR, RF, and SVC. These models were trained using a traditional feature set comprising channel mean, variance, peak, zero-crossings, skew, and kurtosis (Naseer & Hong, 2015). These features were computed per channel separately for HbO and HbR.

The ten Riemannian models are differentiated by variations in feature estimation and classifiers. We assessed a Riemannian-adapted SVC and LR in the tangent space of the Riemannian manifold. The feature estimation techniques explored were as follows:

a) Full kernel matrix estimation: This approach concatenated HbO and HbR signals into a single 124 x 124 kernel matrix. As an SPD matrix in $7^{124 \times 124}$ its Riemannian manifold has dimension $\frac{124 \times 125}{2} = 7750$

b) Block diagonal covariance matrix: We constructed two separate 62×62 covariance matrices for HbO and HbR and placed them on the diagonal, zeroing out cross-terms. This yields a 124×124 block-diagonal SPD matrix whose manifold dimension is the sum of the two 62×62 blocks, i.e. $2 \times \frac{62 \times 63}{2} = 3906$.

c) Block-kernel matrix: Here, the HbO and HbR blocks represented channel kernel functions instead of covariances.

d) Super-kernel block diagonal matrix: In this approach, the HbO and HbR blocks could represent either covariance or kernel matrices.

Kernel matrix analyses reveal complementary HbO and HbR patterns

Only a few studies have utilized channel relationships, such as kernel functions, to classify fNIRS data. However, those who have explored this approach demonstrated that kernels outperformed traditional fNIRS features for classifying mental tasks (Verdière et al., 2018). For kernel matrices to be an effective feature for classification, it is crucial to record the activity from different brain regions as cleanly separated signals (minimizing mixture) and to obtain as many independent brain-activation measurements as possible. This can be optimally achieved by covering a large cortical area. Thus, our montage included 62 optode channels (Fig 3a), spanning the whole left hemisphere from occipital to frontal regions as well as parts of the right hemisphere. We sorted the individual channels of the kernel matrices according to their locations in the 10-20 EEG system in (anterior)-frontal, (fronto)-central, (centro)-parietal, occipital-parietal, temporal, and occipital. Fig 3a shows this sorting for two example kernel matrices (HbO and HbR) and the tasks “Mental Talking” and “Own Paradigm” (all eight tasks for this participant can be found in Supplementary Fig 2). Generally, for the group, kernel values were higher for HbO than HbR (t = 5.08, P = 2.26 x 10^{-3}), likely due to the different signal-to-noise ratios of the two signals. While scaling covariance matrices affects the eigenvalues, it does not affect the eigenvectors.

We tested whether HbO and HbR displayed different co-activation patterns by comparing the distributions of pairwise distances between all possible combinations of two tasks. We found that the average Wasserstein distance was significantly greater than zero, indicating that HbO and HbR pairwise distances contain different information structures (t = 4.68, P = 0.003). Generally, HbO showed a higher average distance between task pairs compared to HbR (t = 3.25, P = 0.017).

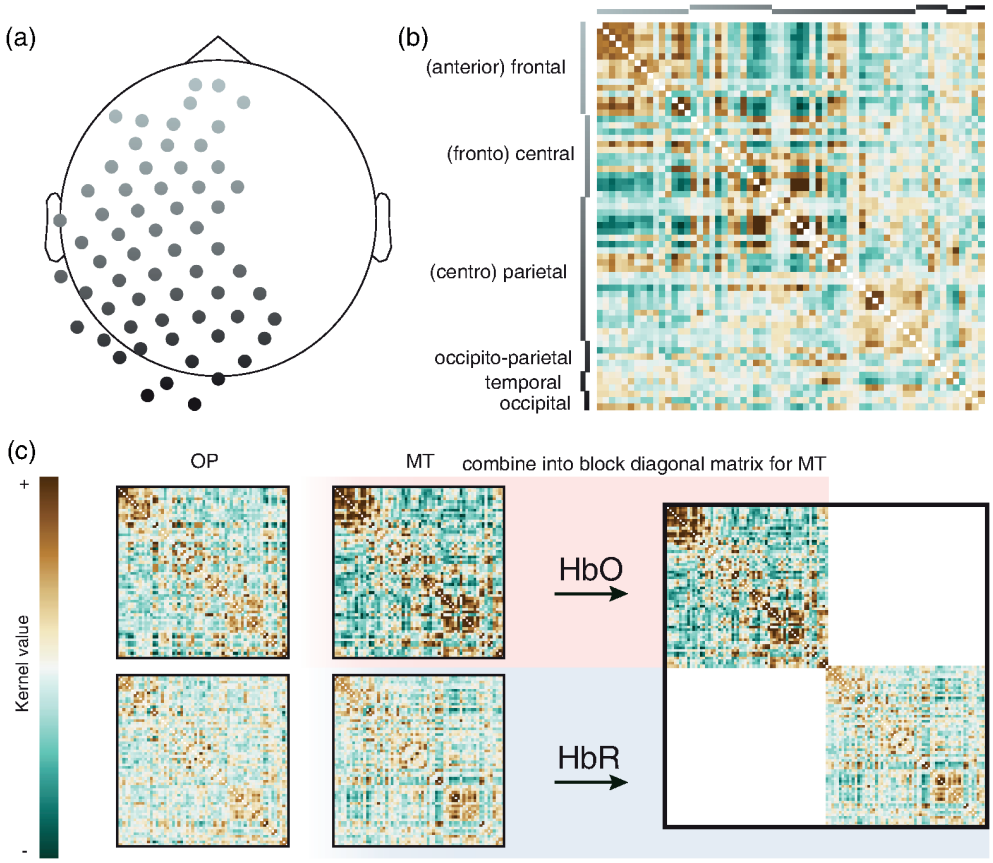

Figure 3. Example average kernel and block diagonal matrices.
(a) FNIRS channel map. **(b)** Example kernel matrix for P01 and task MT. Rows and columns of the matrix are reordered according to the respective EEG locations of the channels. **(c)** Example average kernel matrices from participant P01 for HbO (top row) and HbR (bottom row) for task OP (left column) and task MT (right column). For task MT, the HbO and HbR kernel matrices are stacked into a block diagonal matrix on the right; the same is done for task OP but not illustrated here for space.

The duality of fNIRS data suggests that one could use HbO and HbR as separate features for classification. However, adding more features to classification analyses does not necessarily improve classification accuracy. Riemannian geometry provides an elegant framework for combining different kernel matrices (features) into a single matrix by stacking them along the diagonal, resulting in a block diagonal matrix (Fig 3c). Thus, Riemannian Geometry is particularly well-suited to manage the duality inherent in fNIRS data.

vRiemannian models consistently outperform other models

As the final step of our model fitting process, we classify the previously estimated kernel matrices using either Riemannian-adapted classifiers or traditional machine-learning models. Since our dataset includes eight distinct mental-task conditions, it provides an ideal opportunity to test the potential of co-activation-type classification in distinguishing among multiple tasks. For this 8-choice classification, the traditional models (LDA, LR, RF, SVC) yielded a mean classification accuracy of 42% (SD = 0.120, 95% CI = [0.303, 0.525]) (all $P < 0.01$, permutation tests). In contrast, Riemannian models significantly outperformed traditional models, achieving a mean accuracy of 62% (SD = 0.148, 95% CI = [0.482, 0.755]) (bootstrapped paired *t*-test, $t = 4.81$, $P = 2.98 \times 10^{-3}$, $d = 1.817$; see Fig 4a).

In addition to evaluating the 8-choice classification performance across all tasks, we also explored a binary classification scenario by examining all possible pairwise combinations of the eight tasks, resulting in 28 different binary comparisons. Here, the traditional models achieved a mean accuracy of 78% (SD = 0.073, 95% CI = [0.711, 0.846]) with all classification accuracies above chance level (all $P < 0.01$, permutation tests). The Riemannian models outperformed the traditional models, achieving an average classification accuracy of 95% (SD = 0.020, 95% CI = [0.934, 0.972]) (bootstrapped paired *t*-test, $t = 7.65$, $P = 2.50 \times 10^{-4}$, $d = 2.895$; see Fig 4b).

The “super-kernel SVC" model performed best in both the 8-choice and 2-choice comparisons (65% (SD = 0.138, 95% CI = [0.519, 0.775]) and 96% (SD = 0.112, [0.944, 0.975]) accuracy, respectively). Overall, models utilizing block diagonal matrices (n = 8) as features outperformed those based on full-kernel matrices (models: C-SVC, C-TLR) (n = 2) for both 2-choice comparison ($t = 7.57$, $P = 2.77 \times 10^{-4}$, mean difference = 0.075, SD = 0.026, 95% CI = [0.051, 0.099], $d = 2.86$) and 8-choice comparison ($t = 6.34$, $P = 7.23 \times 10^{-4}$, mean difference = 0.104, SD = 0.043, 95% CI = [0.064, 0.144], $d = 2.39$).

To quantify the improvement, we calculated a relative improvement score comparing the performance of traditional models to Riemannian models. This analysis revealed that the improvement was more pronounced in the 2-choice comparisons than in the 8-choice comparisons for all participants ($t = -8.37$, $P = 1.58 \times 10^{-4}$, mean difference = −0.431, SD = 0.136, 95% CI = [−0.557, −0.305], $d = -3.16$, see Fig 4c).

Using our best-performing model, we analyzed the average confusion matrix for the 8-choice classification (Fig 4d). The results revealed that "Mental Talking" (MT) had the highest accuracy in predicted versus true labels, with 70% correctly classified. This was followed closely by "Mental Calculation" (68%), "Own Paradigm" (68%), "Spatial-Navigation Imagery" (68%), "Tennis Imagery" (62%), "Mental Singing" (58%), and "Mental Drawing" (58%). Interestingly, certain tasks showed complete separability; for example, "Mental Talking," the most distinguishable task, was never mistaken for "Own Paradigm" or "Tennis Imagery." Similarly, "Mental Calculation" was never confused with "Mental Drawing," "Mental Rotation" was never confused with "Mental Singing", and "Own Paradigm" was never confused with "Spatial-Navigation Imagery". These results suggest that certain task pairs may generate more orthogonal co-activation patterns than others, which can be effectively captured by our proposed methods. However, the classification accuracy for each task did not show a significant correlation with participants' ratings of ease and pleasantness (all $P > 0.05$, Spearman rank correlation; see Supplementary Fig 1).

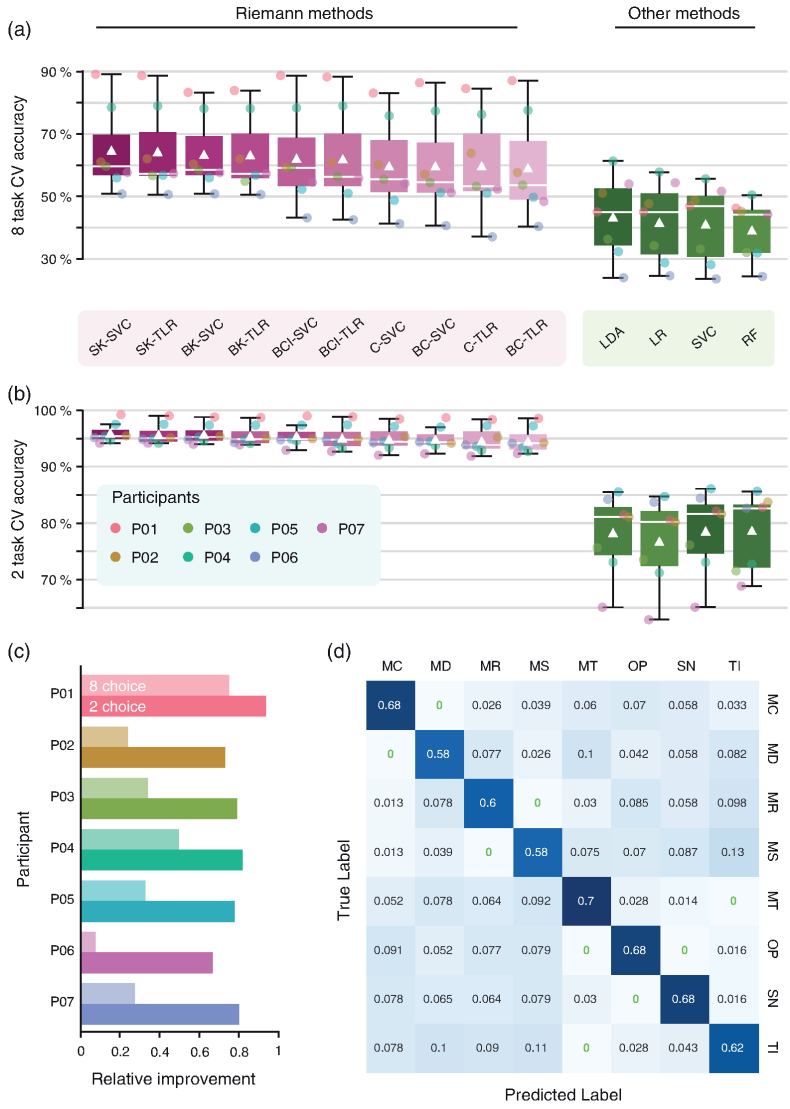

Figure 4. Classification results.

Classification results of ten different Riemann-based methods (pink) and four traditional methods (green). **(a)** The mean accuracies for the one multi-class classification of all combinations of eight tasks. The box plot shows the median for each model as a white bar and the mean as a white triangle. Outliers are depicted as circles. **(b)** same as a, but for the 28 binary classifications of all combinations of two tasks (n = 7 participants in both cases). **(c)** Relative improvements of average accuracies from traditional to Riemann methods per participant. Lighter colors indicate relative improvements for the 8-choice paradigm and darker colors indicate relative improvements for all possible 2-choice combinations. d Mean confusion matrix for the 8-choice comparison, using the best-performing model (US_SVC). Cells for tasks that were never confused with other tasks are marked in green. Abbreviations: SK: super-kernel, BK: block-kernel, C: covariance, BCI: block covariance individual, SVC: (Riemannian) support vector classifier, TLR: tangent-space logistic regression, SVC: support vector classifier, RF: random forest, LDA: linear discriminant analysis, LR: logistic regression.

Discussion

In this study, we demonstrate that classification models based on Riemannian Geometry offer a powerful and robust method for classifying brain states with fNIRS. We demonstrate that Riemannian models consistently outperform traditional machine-learning models in both binary and multi-class classification scenarios (illustrated in Fig 4). Additionally, our study underscores the benefit of leveraging the complementary information inherent to HbO and HbR signals captured by fNIRS. Stacking HbO and HbR signals in block diagonal matrices, as opposed to full covariance matrices, significantly enhances classification performance. Specifically, the combination of a block diagonal matrix as a feature together with a Riemannian Support Vector Classifier (SVC) achieved the highest accuracy among all models in both the 8-choice and 2-choice classification (see Fig 4), with an average of 65% and 96% cross-validation accuracy, respectively.

The benefits of combining Riemannian Geometry with a block diagonal approach

Riemannian models consistently outperformed traditional models such as LR, SVC, LDA, and RF (see Fig 4). This difference in performance likely stems from the distinct features each approach utilizes for classification. While traditional models rely on statistical features computed from individual channels (Herff et al., 2013; Yger et al., 2015), Riemannian models harness information from channel variances and interactions (Barachant et al., 2010). Since brain processes occur within interconnected networks, capturing both linear and non-linear channel interactions can substantially improve classification accuracy (Gemignani et al., 2018). Unlike other statistical features like mean, variance, peak, zero-crossings, skew, and kurtosis, kernel matrices which serve as features in Riemannian approaches, can capture these interactions. For these kernel matrices, non-linear channel kernels can offer an improvement over regular covariances (Wang et al., 2015).

Among the Riemannian models, those employing block diagonal matrices as features outperformed models using full kernel matrices at the group level (see Fig 4). The block diagonal approach allows for the concatenation of multiple SPD matrices, each estimated using potentially different kernel functions and regularization strengths. These models have a larger hyperparameter set, and this additional complexity allows these models to be more finely tailored to the specific characteristics of the two chromophores. The increased flexibility in hyperparameter tuning likely contributed to the superior performance compared to full kernel models. Given that in our data there appears to be minimal meaningful interaction between the HbO and HbR signals, estimating a full kernel matrix may unnecessarily inflate manifold dimensionality without substantially improving classification performance. By setting the interactions between HbO and HbR to zero, block diagonal matrices eliminate these redundancies, resulting in a manifold with lower dimensionality. This submanifold, corresponding to the block diagonal matrix, is therefore more compact and likely better suited

to capture the essential structure of the data. The primary advantage of block diagonal matrices lies in independently optimizing kernel selection and regularization parameters for HbO and HbR, maximizing their distinct informational contributions.

The best-performing model in our study was a super kernel Riemannian SVC (SK-SVC, see Fig 4). As a powerful linear classifier, the SVC is known for its robust generalization capabilities. When data are not linearly separable in their native space, SVC can map them to a higher-dimensional space where linear separation becomes possible, making it highly efficient in capturing potential non-linear interactions. Although SVCs traditionally operate with Euclidean feature vectors, they can be extended to incorporate a Riemannian kernel, allowing for kernel matrices to be used as input features (Keshmiri et al., 2019). In this study, we compared a regular SVC with traditional features to the Riemannian SVC with (block) kernel matrices as features (see Fig 4). The Riemannian SVC, based on the block diagonal approach with separate HbO and HbR tuning, stands out as the most effective method for classifying fNIRS mental imagery data in our study.

Integration of Riemannian Geometry in fNIRS-based clinical applications

A rapidly growing field of fNIRS applications is brain-computer interfacing for motor-independent communication and control, where voluntarily generated brain activations/signals are used to encode computer commands. In these systems, the ability to accurately classify distinct patterns of brain activity is crucial for enabling effective communication and control. As fNIRS-based brain-computer interfaces (BCIs) continue to evolve (Paulmurugan et al., 2021), integrating advanced analytical techniques like Riemannian Geometry has the potential to significantly enhance the precision and reliability of these systems. Similarly, Riemannian Geometry has successfully been used to advance EEG-based BCIs (Barachant et al., 2013) and extending this approach to fNIRS represents a natural progression. By using mental-imagery tasks, crucial for BCI-applications, we show that applying Riemannian Geometry to fNIRS data has great potential for the BCI field. Most notable is the substantial increase in classification accuracy achievable using Riemannian models compared to traditional classifiers (see Fig 4a). Additionally, Riemannian models are computationally efficient and typically require less training data compared to deep-learning models. Lastly, tools like the pyriemann package simplify the implementation of Riemannian models, making them accessible to researchers across various disciplines.

In addition to the general practical benefits of Riemannian models, the block-diagonal approach suggested in this work specifically adds value to real-world fNIRS-based applications. A key advantage of the block-diagonal Riemannian framework in BCI and clinical contexts lies in its ability to allow for different kernel estimations and regularizations for HbO and HbR covariance patterns while substantially reducing model dimensionality relative to full covariance approaches. This reduction in dimensionality in turn can accelerate classifier training and lowers decoding latency. Thus, the block-diagonal Riemannian approach offers a balanced trade-off between representational richness and computational efficiency, strengthening its translational

relevance for clinical BCI implementations. These advantages could make fNIRS, with its unique HbO and HbR signals, a viable alternative to other BCI input modalities, like fMRI and EEG.

Currently, EEG remains the predominant method in real-world BCI applications (Peksa & Mamchur, 2023), with several established classification methods based on Riemannian Geometry (Barachant & Bonnet, 2011; Congedo et al., 2017; Delgado et al., 2020; Gaur et al., 2018; Kumar et al., 2017). Our study demonstrates that Riemannian Geometry may also be well-suited for fNIRS-based BCIs. Thus, the combination of EEG and fNIRS in hybrid BCIs represents a promising future direction (Maher et al., 2023). Future studies could identify co-activation patterns between these two modalities, potentially enhancing BCI performance by leveraging the complementary information from both techniques within a unified Riemannian Geometry framework. The integration of EEG with fNIRS might proceed according to the integration of HbO and HbR exemplified here, namely by means of block diagonal matrices. This integration could further significantly advance BCI technology, offering more robust and versatile systems in which the unique advantages of fNIRS and EEG can optimally be exploited.

For practical BCI applications, it may be advantageous to focus on a smaller subset of tasks, selecting those that are most orthogonal to each other to enhance classification accuracy. Our dataset, which includes eight mental tasks, offers a unique opportunity to identify which tasks are most distinguishable using Riemannian classification (Fig 4). We identified "Mental Talking" as the most easily distinguishable task using our analysis approach, likely due to its intuitive nature and close relation to natural thought processes. Interestingly, participant ratings supported this, with “Mental Talking” ranking second highest for "easiness" and third for "pleasantness" (see Supplementary Fig S1).

Next to the mentioned BCI applications, fNIRS-based classifications using Riemannian geometry could also have significant implications for other clinical applications considering its accuracy, computational efficiency, and accessibility. As an example, ongoing research in our lab applies the novel method to further advance earlier attempts for the brain-based diagnosis of remaining awareness in patients with disorders of consciousness (Owen et al., 2006). In these clinical populations, the hemodynamic response often deviates significantly from typical patterns, potentially differing substantially between HbO and HbR signals. For example, in mental imagery, the signal time course greatly depends on how an individual subject performs the task—whether they begin and end immediately or act with delay. This highlights why our approach, especially extensive hyperparameter tuning and the use of block-diagonal matrices, is particularly suitable for BCIs and other clinical applications.

Limitations and future directions

Despite the promising results, several limitations must be considered. The study involved only seven participants who were students and staff members at Maastricht University, which limits

the generalizability of the findings. Future efforts to assess the generalizability of our findings should include a larger sample and stratify across demographic variables (e.g. age and sex) and physical characteristics (e.g. hair density and type, and skin pigmentation) and include these variables in the analysis to assess their selective effects on our classification approach (Yücel et al., 2024). Also, the study obtained per participant only 96 trials. This modest amount of data has likely resulted in a relatively low sensitivity. This has however not prevented clear and significant results, supporting the possibility for using this approach in clinical practice with typically restricted recording conditions. Nevertheless, future studies with more data per participant might investigate this aspect in more detail. However, clinical research, particularly in BCI applications, often prioritizes individualized analyses, reducing the necessity for broader population generalization. Note also, that despite the limited sample size and the correspondingly low sensitivity, our findings consistently demonstrated that Riemannian models significantly outperformed traditional models at the group level.

However, the observed "complementarity" between HbO and HbR signals that potentially contributed to the superior performance of Riemannian models should be interpreted with caution. We have observed that HbO and HbR signals yield significantly different Riemannian distance distributions, which suggests that each chromophore contributes distinct information to our classifiers. Yet, the canonical hemodynamic response inherently links HbO and HbR in an inverse fashion, and without short-separation channel regression, extracerebral factors (e.g., scalp blood flow, cardiac and respiratory oscillations) may differentially influence each measure. Consequently, the enhanced decoding accuracy obtained by integrating both HbO and HbR within a block-diagonal Riemannian framework likely reflects a combination of true cortical dynamics and divergent artifact sensitivities, rather than two wholly independent neural signals. Future work employing short-distance channels, advanced physiological regression methods, and direct neural validation (e.g., simultaneous EEG) will be critical for disentangling cortical contributions from systemic noise and for confirming the neuronal underpinnings of any observed HbO/HbR divergences.

While the present study was principally designed to assess the feasibility of applying Riemannian geometry-based classifiers to fNIRS data rather than to perform a direct comparison of individual mental-task paradigms, our optode montage with a left-hemispheric focus may have conferred uneven signal-to-noise advantages across tasks. Tasks that elicit bilateral activation (e.g., aspects of spatial navigation or complex imagery) might have been underrepresented in our decoding performance metrics. Although classification accuracies were highest for these conditions, we acknowledge that task-specific optode arrangements or fully bilateral layouts might further enhance decoding robustness. Future investigations should therefore explore tailored or bilateral sampling schemes to maximize sensitivity for each cognitive paradigm, thereby refining both the generalizability and the task-specific efficacy of Riemannian classifiers in fNIRS.

Another limitation is the absence of a dedicated train-test split, which might lead to an overestimation of classification accuracy. Rather, we use cross-validation accuracy, which can inflate real-world performance. This could in principle be mitigated by a train-test split of the available data, which was however not feasible due to the small number of trials per subject. Even though cross-validation does not provide an absolute measure of generalization it allowed for a fair comparison between models, as the same validation protocol and data partitions were applied uniformly across all approaches. Future work may investigate dimensionality reduction strategies (e.g., Riemannian tangent space mapping with feature selection or regularization techniques) and incorporate larger datasets to enable dedicated train-test splits.

Additionally, Riemannian models required more extensive hyperparameter tuning than traditional models. Yet, this extended tuning process allowed us to optimize model performance and reach superior performance. Enhanced model sensitivity enables the detection of subtle neural effects, making this approach valuable for advancing our understanding of brain function. While accuracy is similarly important for BCI applications, time is often a limiting factor in applied fNIRS-BCI settings. BCIs using fNIRS generally work on longer time scales due to the characteristic time delay of the hemodynamic response. Additionally, conducting a large hyperparameter tuning might lead to long fitting times. Thus, it may be necessary to reduce the parameter space in BCI and other clinical contexts to achieve an optimal balance between model-fitting time and accuracy. Alternatively, a Riemannian model with a single kernel matrix and a subset of channels could offer a solution that requires less powerful hardware while still surpassing traditional approaches. Future research might investigate the integration of Riemannian Geometry with deep-learning techniques to further improve classification performance, a strategy that has shown promise in EEG (Paillard et al., 2024).

Furthermore, it should be noted that we did not perform a short-distance channel correction to denoise the fNIRS data, which is currently considered best-practice. However, at the time of data collection, our fNIRS Scout 16x24 system did not support short distance channels which is why we were unable to perform this correction. To mitigate global physiological interference, we instead applied a PCA-based global component regression. This PCA approach is a viable alternative, when short-distance channel measurements are unavailable, to remove heartbeat, respiratory noise, and skin-blood-flow (Virtanen et al., 2009; X. Zhang et al., 2016; Y. Zhang et al., 2005). Nonetheless, our classifiers achieved robust performance across tasks, indicating that task-relevant hemodynamic signals were preserved despite residual systemic noise. Future studies should incorporate short-distance channels to more precisely disentangle cortical from extracerebral contributions.

Due to the short trial durations, the data per trial are also relatively short. We computed covariance matrices over fixed-duration windows to maintain consistency across trials and conditions. Longer recordings per trial would indeed provide more stable covariance estimates

and potentially improve stationarity assumptions, but this would come at the cost of increased trial durations and participant fatigue. Our choice reflects a balance between signal reliability and experimental feasibility.

In the current study, "Mental Talking" was the most distinguishable task out of all eight tasks. Of note is that mental talking may have elicited unconscious movements of the tongue or jaw. These subliminal movements could have contributed to the classification by increasing motor cortical activation. However, previous studies have used “Mental Talking” successfully for fNIRS-based classification in a movement-disabled patient (Gallegos-Ayala et al., 2014), further supporting the distinguishability of the task. Future studies may measure electromyography to detect potentially confounding muscle activity. Note that any such subliminal movements would have been equally available to all classification methods and cannot explain the observed differences between methods.

To validate the utility of a Riemannian Geometry-based approach, the models employed in our rather controlled setting should be tested across diverse settings, from basic research to applied contexts, to determine which models are most effective in various scenarios. Our fNIRS recordings were acquired in a darkened, sound-attenuated laboratory setting chosen to minimize light fluctuations, acoustic distractions, and motion-induced artifacts. By reducing these sources of physiological and optical noise, we were able to obtain a clear benchmark of decoding performance and rigorously compare Riemannian-geometry classifiers under standardized conditions. However, typical BCI deployments (i.e. clinical wards, rehabilitation centers, or home environments) naturally contain environmentally induced noise which can degrade signal quality. Therefore, future work should extend our findings by systematically evaluating classifier robustness in more ecologically valid contexts. Such studies should investigate adaptive preprocessing strategies (e.g., real-time motion correction, dynamic light compensation), optimize optode layouts for mobile applications, and retrain models on data collected in situ to ensure reliable performance outside the laboratory.

Acknowledgements

This work was supported by The Netherlands Organization for Scientific Research (NWO; Vidi-Grant VI.Vidi.191.210 to BS) and by The Luik 3 grant for joint research between Cognitive Neuroscience and Knowledge Engineering Departments, Maastricht University, on advanced brain–robot interfaces, (Grant No. 2015-9999) (to BS). GPT-4.5 was used for grammar checking. An earlier version of this manuscript has been posted as a preprint: https://doi.org/10.1101/2024.09.06.611347.

Supplementary Material

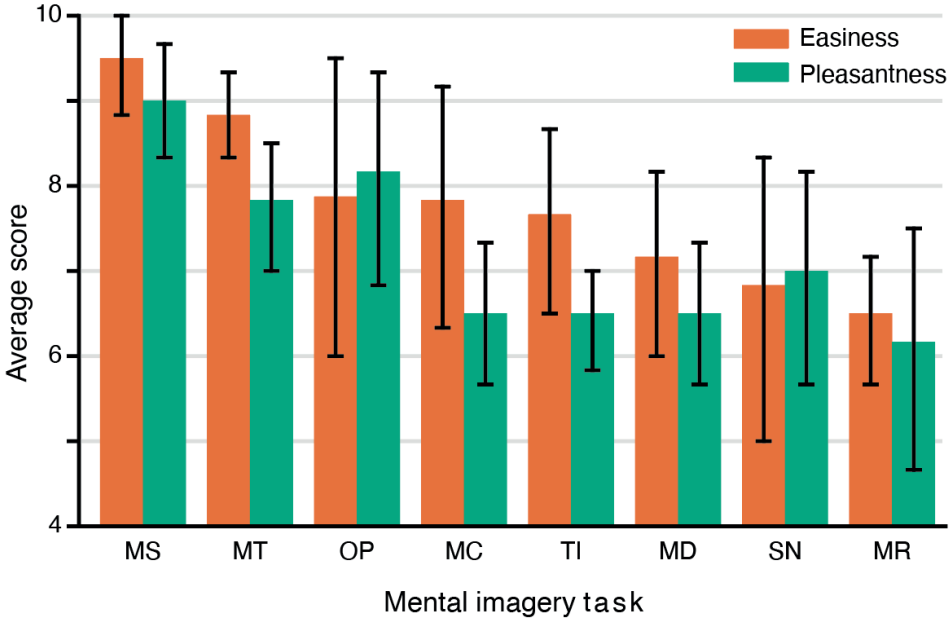

Supplementary Figure S1. Mental task easiness and pleasantness ratings.

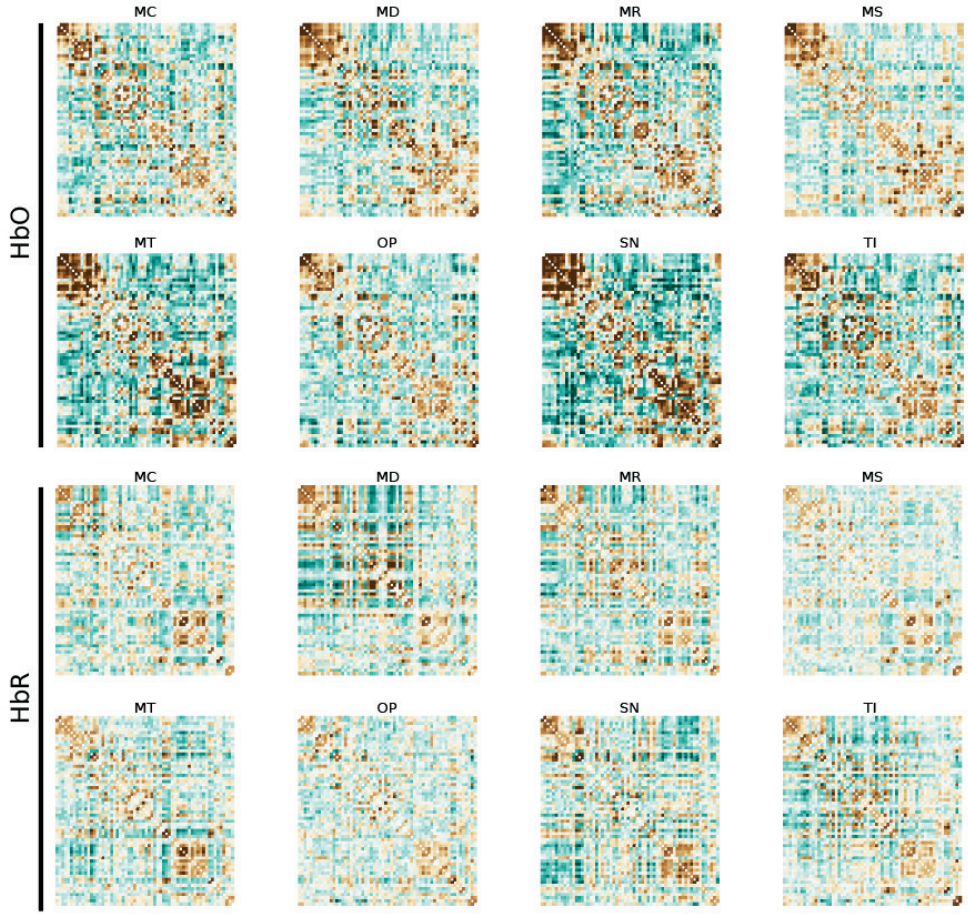

Supplementary Figure S2. Example covariance matrices.

Supplementary Table 1. Participant Characteristics.

Participant	Sex	Age	Head (Cap) size	Handedness	Individually chosen mental task
P01	female	38	56 (56)	right	Visualize horse-racing
P02	female	23	56 (56)	right	Karate movement imagery
P03	male	54	59 (60)	right	Mental planing
P04	male	35	57 (58)	left	Mantra meditation
P05	male	28	58 (58)	right	Mental sex
P06	male	26	56 (56)	right	Mental motorcycling
P07	female	21	56 (56)	right	Mental knitting

Supplementary Table 2. Hyperparameters.

Model Variant	Metrics / Estimators	Shrinkages	C values	SVC Metric Options
SuperKernel SVC	rbf, polynomial, laplacian, cosine, poly, corr, cov, scm, lwf, oas, hub, sch, tyl	0, 0.01, 0.1	0.1, 1, 10, 100	euclid, riemann, logeuclid
SuperKernel TLR	rbf, polynomial, laplacian, cosine, poly, corr, cov, scm, lwf, oas, hub, sch, tyl	0, 0.01, 0.1	0.1, 1, 10, 100	–
BlockKernels SVC	rbf, polynomial, laplacian, cosine, poly	0, 0.01, 0.1	0.1, 1, 10, 100	euclid, riemann, logeuclid
BlockKernels TLR	rbf, polynomial, laplacian, cosine, poly	0, 0.01, 0.1	0.1, 1, 10, 100	–
BlockCov SVC	corr, cov, scm, lwf, oas, hub, sch, tyl	0, 0.01, 0.1	0.1, 1, 10, 100	euclid, riemann, logeuclid
BlockCov TLR	corr, cov, scm, lwf, oas, hub, sch, tyl	0, 0.01, 0.1	0.1, 1, 10, 100	–
BlockCovInd SVC	corr, cov, scm, lwf, oas, hub, sch, tyl	0, 0.01, 0.1	0.1, 1, 10, 100	euclid, riemann, logeuclid
BlockCovInd TLR	corr, cov, scm, lwf, oas, hub, sch, tyl	0, 0.01, 0.1	0.1, 1, 10, 100	–
Cov SVC	corr, cov, scm, lwf, oas, hub, sch, tyl	0, 0.01, 0.1	0.1, 1, 10, 100	euclid, riemann, logeuclid
Cov TLR	corr, cov, scm, lwf, oas, hub, sch, tyl	0, 0.01, 0.1	0.1, 1, 10, 100	–



3

Chapter 3

Sensitive and Specific fNIRS-Based Approach for Awareness Detection in Disorders of Consciousness

This chapter is based on the following publication:

Bastian, L.*, **Näher, T.*,** Vorreuther, A., Lührs, M., Benitez Andonegui, A., Fries, P., Riecke, L., & Sorger, B. (2025). Sensitive and specific fNIRS-based approach for awareness detection in disorders of consciousness: Proof of principle in healthy adults. Neurophotonics, 12(4), 045001.

*These authors contributed equally

DOI: https://doi.org/10.1117/1.NPh.12.4.045001

.

Abstract

Introduction: Disorders of Consciousness (DoC) are characterized by varying levels of arousal and awareness. Due to severe motor impairments often accompanying these conditions, differentiating between minimally conscious state (MCS) and unresponsive wakefulness syndrome (UWS) remains challenging. Accurate diagnosis, however, is critical for optimal treatment and prognosis. Functional Near-Infrared Spectroscopy (fNIRS), due to its portability and non-invasiveness, holds promise for brain-based diagnostics, although current methods lack sufficient sensitivity and specificity.

Methods: This proof-of-concept study introduces a novel, highly sensitive, and specific fNIRS-based diagnostic approach tailored to the individual cognitive state of DoC patients. Nine healthy participants received auditory cues instructing them to either perform individually tailored motor-speech imagery or remain at rest. In separate runs, participants were instructed to be responsive to these cues, to mimic MCS, or unresponsive, to mimic UWS. FNIRS-channel covariance matrices were classified for responsive and unresponsive states as either imagery or rest using a Riemannian geometry-based approach.

Results: Classification between responsive and unresponsive states achieved a sensitivity of 100% and a specificity of 89% across participants. Within the two states, imagery and rest were classified with 83.9% and 55.93% accuracy respectively; the latter result, close to chance level, was expected in the unresponsive state.

Conclusion: This individualized diagnostic approach may have the potential to significantly enhance diagnostic accuracy for DoC. It provides a non-invasive, efficient, and objective assessment, potentially reducing the rate of misdiagnoses rates. The practicality and minimal technical requirements of fNIRS further support future clinical implementation.

Introduction

Disorders of Consciousness (DoC) are clinically defined by changes in arousal and/or awareness (Laureys et al., 2010), classifying these conditions into three distinct states: 1) minimally conscious state (MCS) characterized by full arousal and minimal and inconsistent awareness, 2) unresponsive wakefulness syndrome (UWS) defined by full arousal but no awareness, and 3) coma, marked by an absence of both arousal and awareness (Edlow et al., 2021; Thibaut et al., 2019). While the behavioral diagnosis of coma is unequivocal, differentiating between MCS and UWS remains challenging (Kondziella et al., 2020). Patients in both states exhibit full arousal, but may differ in their level of awareness, which is often difficult to detect. The most common diagnostic methods used in clinical practice, such as the Coma Recovery Scale-Revised (Giacino et al., 2004) and the Glasgow Coma Scale (Teasdale & Jennett, 1974), exclusively rely on overt and controlled motor behavior. However, motor impairments, abnormal sleep-wake cycles, and uncontrolled movements in DoC patients create difficulties in the application of such behavioral assessments. This leads to approximately 40% of MCS patients being inaccurately diagnosed as UWS (Schnakers et al., 2009). Accurate differentiation between UWS and MCS is critically important, as MCS patients generally have a more favorable prognosis (Coleman et al., 2009) and are capable of communicating their needs through brain-computer interfacing (BCI) (Schiff et al., 2024). Furthermore, an incorrect diagnosis of UWS may lead to ethically challenging decisions, such as the termination of life-sustaining treatment, underscoring the urgent need for more accurate diagnostic tools (Andrews et al., 1996; Childs et al., 1993).

In light of these challenges, alternative, brain-based diagnostic approaches have been explored, combining instructions to perform an imagery task and non-invasive functional neuroimaging methods. These methods leverage mental rather than motor responses, with Electroencephalography (EEG) and functional Magnetic Resonance Imaging (fMRI) being the most extensively studies modalities (Bardin et al., 2011, 2012; Claassen et al., 2019; Cruse et al., 2011; Galiotta et al., 2022; Gibson et al., 2014; Luauté et al., 2015; Monti et al., 2010; Owen et al., 2006) (for a recent review see (Schiff et al., 2024)). Despite their promise, both EEG and fMRI have notable limitations. EEG provides a high temporal resolution, yet it suffers from a substantial susceptibility to ocular and muscle-related artifacts. EEG measurements are also less sensitive in DoC patients with traumatic brain injury due to skull deformities (Owen et al., 2006). In contrast, fMRI offers excellent spatial resolution through measuring blood-oxygen-level-dependent (BOLD) signals but lacks portability and bedside applicability due to its stationary nature. Functional Near-Infrared Spectroscopy (fNIRS) emerged as a promising alternative functional imaging modality for detecting awareness (Abdalmalak et al., 2021). Most importantly, fNIRS simultaneously measures changes of both oxygenated (HbO) and deoxygenated (HbR) hemoglobin, which can provide complementary sources of information (Villringer et al., 1993) and therewith potentially more diagnostic certainty. FNIRS is affordable, practical and quieter compared to fMRI, and has reduced susceptibility to artifacts compared to

EEG (Klein et al., 2024). However, it is limited by potential signal interference from hair density and skin pigmentation (Benitez-Andonegui et al., 2021); for a detailed evaluation see Klein et al. (2024).

Despite its significant and unique clinical potential, fNIRS-based awareness detection in DoC patients has been explored in only a small number of studies. Existing research has identified residual awareness through resting-state functional connectivity (Liu et al., 2023; Shu et al., 2023), passive paradigms (Lu et al., 2023; Molteni et al., 2013), or active paradigms (Abdalmalak et al., 2020; Kazazian et al., 2024; Kempny et al., 2016; Molteni et al., 2013; Si et al., 2023). Active paradigms generally yield more robust and reliable results (Abdalmalak et al., 2021) and can subsequently be employed for BCI-mediated communication(Vorreuther et al., 2023). Using an active paradigm in MCS/UWS patients, Abdalmalak and colleagues (2020) achieved 79-83% diagnostic sensitivity and 58-71% specificity, highlighting its great promise for DoC diagnosis (Abdalmalak et al., 2021). However, we believe that experimental designs and procedures need to and can be further improved.

A potential area for improvement is the implementation of the mental-imagery tasks. Existing active paradigms rely on predefined imagery tasks with long encoding times (7-30 s), which may impose unnecessary cognitive demands and overlook individual differences in patient capabilities due their unique patterns of brain damage. To overcome this, we here suggest and test an individualized imagery task, composed of a motor and speech component, with a short encoding time of only 2 s. Importantly, patients could choose the exact motor- and speech-imagery strategy themselves depending on their unique mental capabilities. Another potential limitation of existing paradigms is the analysis method. All the studies mentioned above relied on a canonical hemodynamic response function (HRF) that assumes a uniform response shape. In addition, most of the studies did not consider multi-channel activation patterns which, by definition, provide more, potentially meaningful, information than single-channel or channel-average analyses. To address this, we used a more sensitive covariance-based classification approach that considers co-activation patterns between fNIRS channels without relying on a predefined HRF shape. By reducing cognitive demands and employing a more sensitive analysis approach, our paradigm has the potential to improve diagnostic accuracy and efficiency in clinical practice.

The present study validated the novel diagnostic approach in healthy adults. Per trial, participants received auditory task cues instructing them to either perform an individualized motor-speech IMAGERY task for 2 s or a NO-IMAGERY resting task for 2 s. Furthermore, in separate experimental runs, participants were instructed to either be RESPONSIVE to the trial-wise task cues, thereby mimicking MCS, or to be UNRESPONSIVE, thereby mimicking UWS. We classified IMAGERY vs. NO-IMAGERY conditions in both the RESPONSIVE and UNRESPONSIVE states. We hypothesized that classification accuracy significantly exceeds chance level in the RESPONSIVE but not in the UNRESPONSIVE state. Using a Riemannian-based classification approach, our diagnostic tool achieved high sensitivity and specificity, demonstrating potential for improvements in the differential diagnosis of DoC patients.

Methods

Participants

We recruited 14 healthy adults with age = 27.5 ± 11.4 years (mean ± SD) and normal hearing, including external volunteers, and students and staff members of the Faculty of Psychology and Neuroscience at Maastricht University. Five participants were subsequently excluded from the analysis, for the following reasons: Two participants were excluded due to data corruption, one participant had insufficient fNIRS signal quality due to physical attributes (i.e. hair thickness and skin pigmentation), another participant fell asleep during the task after extensive physical exercise before the fNIRS session, and the fifth participant reported claustrophobic symptoms in the test cabin (see Tab. 1 for participant characteristics). Participants gave written informed consent and received financial compensation for their participation. The local ethics committee approved the experiment (Ethics Review Committee Psychology and Neuroscience [ERCPN]).

Table 1. Participant Demographics.

Participant	Sex	Age (years)	Handedness	Cap Size Used (cm)
P01	F	21	Right	58
P02	F	47	Right	56
P03	F	23	Right	56
P04	F	25	Right	56
P05	F	23	Right	58
P06	F	20	Right	58
P07	F	22	Right	56
P08	F	27	Right	58
P09	M	22	Right	58

Note. F = female; M = male.

General procedure

The experiment consisted of a training session (Fig. 1A) and the fNIRS session. During the training session, participants were familiarized with the mental-imagery task they had to perform in the IMAGERY condition within the RESPONSIVE state. The task consisted of 2 s-long, combined motor-speech imagery. Since the encoding time of 2 s is very short, and some DoC patients may have impairments in either the motor or speech domain, we decided to combine motor and speech components to increase the likelihood to evoke a robust hemodynamic response in as many participants as possible (Holper et al., 2009). To allow for a user-centered approach, participants could choose the specific strategy to perform the task. To guide the participant's choice of strategy, they were asked to think of how they would naturally (within approximately 2 s) draw attention to themselves if they were in the situation of a DoC patient by imagining to, *e.g.*, wave hands/knock and shout "Hello, I am here!". The training session ended with measuring the head size and determining the respective fNIRS-cap size.

One to two days after the training session, participants performed an fNIRS session (Fig. 1B). At the beginning of the session, participants were seated comfortably in a noise-dimmed cabin. The fNIRS cap was placed on the participant's head and all optodes were adjusted to ensure orthogonality to the scalp and proper optode-scalp contact. Participants were instructed to sit still, close their eyes, and avoid any movements during the recordings. The fNIRS session lasted for approximately 1h including participant preparation and data-acquisition (Fig. 1B).

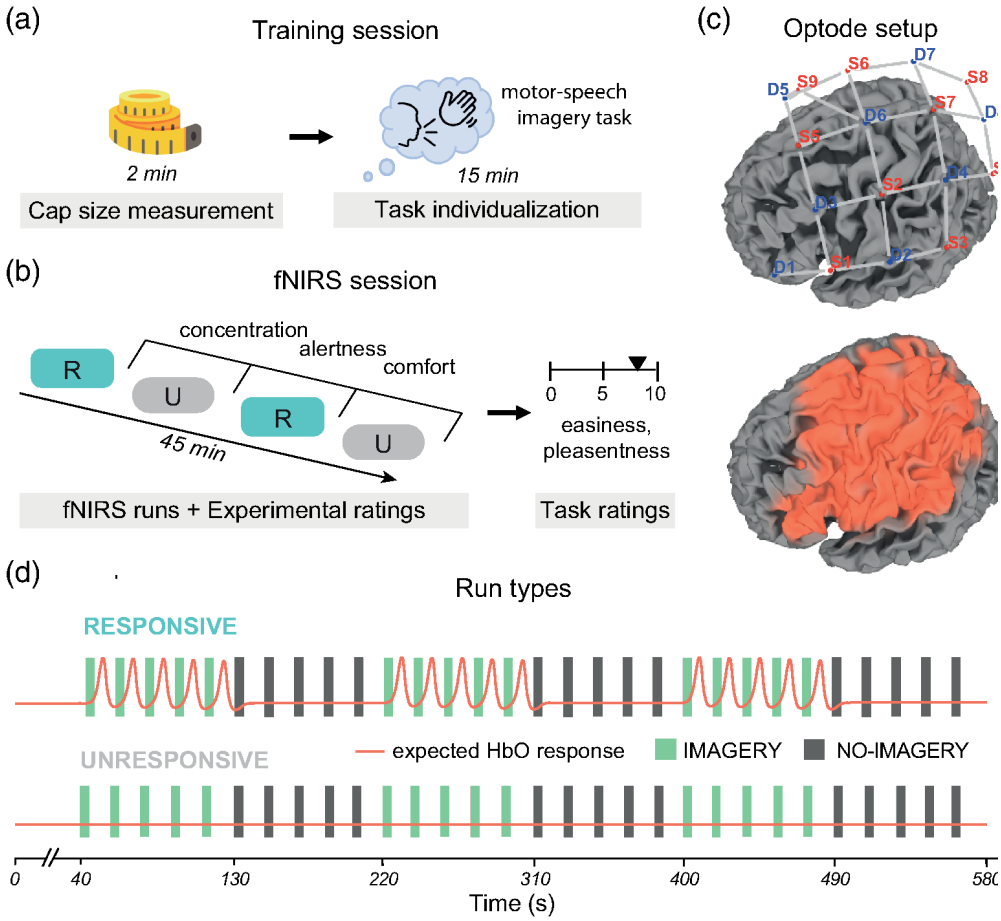

Figure 1. Experimental procedure.

A, Prior to fNIRS-data collection, participants completed a training session. This session included cap-size determination and a 15-min task-individualization procedure in which subjects could define their own motor-speech imagery strategy for the task. **B**, The fNIRS session took approximately 1h and consisted of four (two RESPONSIVE and two UNRESPONSIVE) fNIRS runs lasting 09:40 min each. The runs alternated between RESPONSIVE and UNRESPONSIVE. Half of the participants started with RESPONSIVE, and the other half with UNRESPONSIVE. In between each run, participants rated their level of experienced alertness, concentration, and comfort. After data collection, participants indicated easiness and pleasantness of the mental-imagery task and the resting task. **C**, The optode setup included eight detectors and nine sources over the left hemisphere covering both motor and speech areas. The brain regions marked in orange show the light-intensity map and thereby the regions covered by the fNIRS montage. We also included a short-distance channel (S9-D5). **D**, Schematic representation of the blocked event-related design for RESPONSIVE and UNRESPONSIVE runs. IMAGERY and NO-IMAGERY trials were performed in blocks of five. Each trial consisted of a 1-s auditory cue of high or low frequency and a 2-s task time. Between trials, there were 15 s resting periods. The allocation of the two sound frequencies (high and low) to prompt the two task conditions (IMAGERY or NO-IMAGERY) was balanced across the four fNIRS runs.

Imagery paradigm

During the fNIRS session, participants performed four runs, each lasting 09:40 min. Each run started with a 40-s baseline period after which participants performed six blocks of 5 trials each, alternating between IMAGERY and NO-IMAGERY blocks. The order of blocks was counterbalanced across runs. Each trial lasted 3 s including an auditory cue of high (1000 Hz) or low frequency (300 Hz) of 1 s, a task period of 2 s, followed by a resting period of 15 s. Out of the four runs, two runs were RESPONSIVE, while the other two were UNRESPONSIVE (Fig. 1B). RESPONSIVE and UNRESPONSIVE runs were performed in alternation; half of the participants started with RESPONSIVE and the other half with UNRESPONSIVE. We used two separate RESPONSIVE and UNRESPONSIVE runs to account for within-session variability of fNIRS data quality and alertness, i.e. to reduce potential fatigue and attentional drift effects within a single long block. The RESPONSIVE state was used to simulate an MCS patient, who is able to follow a command by willful modulation of brain activity (Fig. 1D). For IMAGERY trials, participants performed the mental task in response to one type of auditory cue (high- or low-frequency sound). For NO-IMAGERY trials, they did not perform the mental-imagery task indicated by the other sound. The UNRESPONSIVE state was used to simulate a UWS patient, who cannot show mental command following. Here, participants were instructed to never perform the mental-imagery task, neither in IMAGERY nor in NO-IMAGERY trials, which were again indicated by two different sound frequencies (Fig. 1D). The order of auditory cues was counterbalanced across participants for IMAGERY and NO-IMAGERY trials.

Data acquisition

A NIRScout-816 system (NIRx Medizintechnik GmbH, Berlin, Germany) was used to measure the hemodynamic response, using nine source and eight detector optodes (Fig. 1C). Sources were positioned according to the international 10-20 EEG system on FC5 (1), C3 (2), CP5 (3), P3 (4), FC1(5), Cz (6), CP1 (7), Pz (8) and at the midsagittal dimension between Cz and FCz (9). Detectors were positioned on F5 (1), C5 (2), FC3 (3), CP3 (4), FCz (5), C1 (6), CPz (7) and P1 (8). Source 9 and detector 5 formed a short-distance channel with a 1-cm optode distance to obtain information on physiological noise (i.e., Mayer waves [≈0.1Hz]) within the superficial layers excluding brain tissue and affecting specifically HbO time courses. Information obtained from this channel was used later to regress out systemic noise from the 3-cm normal-distance channels. The placement of the entire cap took on average approximately 5 min. Placed optodes covered a region above the left-hemispheric fronto-parietal cortex where we expected motor and speech task-related modulation of brain activation. Data were recorded using NIRStar 15.2 (*NIRx Medizintechnik GmbH, Berlin, Germany*) at 6.94 Hz.

Subjective ratings

After each run, participants rated their level of experienced concentration, alertness, and comfort according to a Likert scale ranging from 1 (extremely unconcentrated/tired/uncomfortable) to 10 (extremely concentrated/alert/comfortable). Upon completion of the

experiment, participants rated the general easiness and pleasantness of the mental imagery task and resting task on a Likert scale ranging from 1 (extremely difficult/unpleasant) to 10 (extremely easy/pleasant).

Data analysis

Data preprocessing

The fNIRS data were preprocessed with the MNE library (Gramfort et al., 2013) in Python 3.10. Initially, the raw data were visually inspected during data acquisition for salient signal increases after trial onset. We again inspected the data offline prior to the analysis using pulse oscillations as a marker for signal quality. Based on the inspection, no trials were removed. Then, the raw signals were converted into optical densities. After this transformation, we used temporal derivative distribution repair (Fishburn et al., 2019) to rectify baseline shifts and eliminate spike artifacts. Short distance-channel regression was applied using the MNE implementation (Fabbri et al., 2004; Saager & Berger, 2005; Scholkmann et al., 2014) to remove Mayer waves and other extracerebral noise. Next, optical densities were converted to hemoglobin-concentration changes applying the modified Beer-Lambert law with a partial path length factor of 6. The data were then bandpass-filtered using a second-order Butterworth filter (low-pass and high-pass cut-offs at 0.4Hz and 0.1Hz, respectively) and segmented into epochs. Each epoch was defined starting from the auditory cue onset and extending to 12 s post-onset. Linear trends were removed from each epoch.

Classification

Using a novel approach, the analysis aimed to classify participants as RESPONSIVE, mimicking MCS, or UNRESPONSIVE, mimicking UWS. Typically, classifiers are trained on a dedicated training set and then used to make predictions for new, unseen data. In contrast, our approach utilized the successful discrimination between IMAGERY and NO-IMAGERY as an indicator of the participant's compliance with the task instructions. The logic of this approach is as follows: If DoC patients can adhere to instructions (i.e., they are, for example, in an MCS), they should be able to modulate their brain activity by performing a mental-imagery task. When trained on such data, this modulation implies that a classifier can distinguish between IMAGERY (when the patient engages in the imagery task) and NO-IMAGERY (when the patient continues to rest) trials. Conversely, if a DoC patient is unresponsive (i.e., is in UWS) and thus unable to follow the task instruction, their brain activity should not exhibit such modulation. Consequently, in such cases, a classifier should fail to differentiate between IMAGERY and NO-IMAGERY trials, resulting in classification accuracies close to chance level. Each participant mimicked both an MCS and a UWS patient by being RESPONSIVE in two runs and UNRESPONSIVE in the other two runs.

We combined trials from corresponding runs for classification, resulting in 60 trials per responsiveness condition in each participant (Fig. 2A). We then computed the channel covariance matrices for the two combined runs (cf. *Covariance Estimation and Classifier*). We

conducted a hyperparameter search to determine optimal classifier model parameters. For each hyperparameter set, we obtained classification accuracies by a stratified 5-fold cross-validation. Based on these accuracies, we selected the hyperparameter set that yielded the highest classification accuracy and retrained the respective model (Fig. 2C). Hyperparameter tuning was performed separately for each participant and signal on that participant's own data, and because tuning and evaluation were not nested, the reported accuracies may be slightly optimistic. All analyses were performed separately for HbO, HbR, and combined (i.e., HbO&HbR) data. Note that we used full covariance matrices as features for our classifier, resulting in the dimensionality of combined (HbO&HbR) being larger than that of just HbO or HbR. The dimensionality of the manifold that is used for classification is given by the number of independent elements in the covariance matrix:

$$dim\left(M\right) = \frac{n(n+1)}{2}$$

where *M* is the manifold of $n \times n$ symmetric positive definite matrices (e.g. covariance matrices). Additionally, the combined matrix thus contains the covariances between HbO and HbR.

Covariance estimation and classifier

We used a classification based on the covariance matrices computed over all channels (Fig. 2B). Applying covariance-based classification is a recent development in the fNIRS data analysis field that has shown promising results as a sensitive method for detecting small differences in mental brain states (Näher et al., 2025). Using the covariances between channel pairs has the advantage that no prior channel selection is required and no canonical HRF needs to be assumed. Here, we estimated a special form of the covariance matrix that is specifically useful for the analysis of event-related potentials (Barachant et al., 2013; Barachant & Congedo, 2014). Briefly, a prototyped response *P* is obtained by computing the average signal separately over IMAGERY and NO-IMAGERY trials in the training set:

$$P = \frac{1}{m}\sum_{i=1}^{m} X_i$$

Where *m* is the number of respective trials, and X_i represents the time series data of trial *i*. The prototyped response P is then concatenated with X_i

$$\tilde{X}_i = \begin{bmatrix} P \\ X_i \end{bmatrix}$$

Conceptually, this is equivalent to computing the average response over trials for each channel and treating these averages as new channels. The resulting concatenated trials $\tilde{X}_i$ are then used for covariance estimation. Adding the average, event-related responses of IMAGERY and NO-IMAGERY allows us to take the spatial structure of the signal into account by including the covariances of the individual trials with the prototypes' responses represented by *P*. Note that for HbO&HbR, each channel was represented by both HbO and HbR data, effectively doubling the channel count.

To reduce the dimensionality of covariance matrices and speed up computation during model fitting, we applied the xDAWN algorithm, a spatial filtering technique commonly used in EEG analyses, to boost the signal-to-noise ratio in event-related potentials (Rivet et al., 2009, 2011). Before concatenation, the xDAWN filters were applied to both the prototypes P and the individual trial data X_i. For the final covariance estimation, we used the Regularized Oracle Approximating Shrinkage (OAS) method (Chen et al., 2009). Regularized covariance estimation provides more numerical stability when covariances are estimated on small samples as is the case here. Additionally, we compared different regularizations and found that OAS demonstrated the highest numerical class distinctiveness in differentiating IMAGERY and NO-IMAGERY trials in the RESPONSIVE condition (see *Class distinctiveness* section for more details).

Covariance matrices computed from linearly independent data are inherently symmetric and positive definite, meaning they belong to the space of symmetric positive definite (SPD) matrices. This space forms a Riemannian manifold—a curved space that locally approximates flat Euclidean geometry through its tangent spaces. Since each point on the manifold has an associated tangent space that "touches" the manifold, we can perform well-defined geometric operations in these local coordinates. In our approach, we exploit this structure by projecting the spatially filtered and regularized covariance matrices onto their corresponding tangent spaces. This projection maps each covariance matrix onto a Euclidean vector whose dimensionality depends on the number of spatially filtered channels, effectively preserving the local structure of the original manifold. We employed a Riemannian metric both for estimating the reference matrix and for performing the tangent-space mapping, following the methodology described by Barachant and colleagues (2012). The reference matrix, typically computed as the Riemannian mean of the covariance matrices, serves as the common anchor point from which the tangent-space mapping is derived, preserving the intrinsic geometry of the SPD manifold. Finally, these Euclidean-transformed vectors were classified using binary logistic regression to distinguish between IMAGERY and NO-IMAGERY. The hyperparameter search for this pipeline included the number of spatial filters for the xDAWN algorithm (2, 4, or 6 filters) and the regularization term *C* for the logistic regression (1, 10, 100, or 1000).

Class distinctiveness

We compared the class distinctiveness on the manifold of covariance matrices for the two conditions IMAGERY vs. NO-IMAGERY. We followed the definition from (Izzuddin et al., 2021):

$$\mathrm{classDis}(A,B) = \frac{d(\bar{X}^A, \bar{X}^B)}{\frac{1}{2}\left(\sigma_{X_A} + \sigma_{X_B}\right)}$$

where $\bar{X}^A$ and $\bar{X}^B$ are the centers of IMAGERY and NO-IMAGERY, i.e., the mean of matrices from each class, and $\sigma_{X_A}\,\sigma_{X_B}$ are the class dispersions, i.e., the mean of distances between matrices from class K and their center of class:

$$\sigma^p_{X_K} = \frac{1}{m}\sum_{i=1}^{m} d(X_i, \bar{X}^K)^p$$

We found OAS estimation of covariances resulted in the highest numerical class distinctiveness. OAS regularization was then chosen for all subsequent covariance estimations.

Thresholding and permutation tests

To classify each participant dataset as RESPONSIVE or UNRESPONSIVE, we used permutation tests (Combrisson & Jerbi, 2015), which involved randomly shuffling the IMAGERY and NO-IMAGERY trial labels 250 times to obtain a null distribution and comparing it to the observed cross-validation accuracy to obtain an empirical *P*-value. We applied this test independently to HbO, HbR, and the combined (HbO&HbR) data (Fig. 2D). To obtain a final prediction for each state, we took the minimum *P*-value of the three separate analyses (HbO, HbR, and HbO&HbR). A dataset was classified as RESPONSIVE if the *P*-value was below the threshold of $P = 0.05$ (Fig. 2E). Such a result would suggest that the observed cross-validation accuracy significantly exceeded what would be expected by chance, indicating the condition as RESPONSIVE.

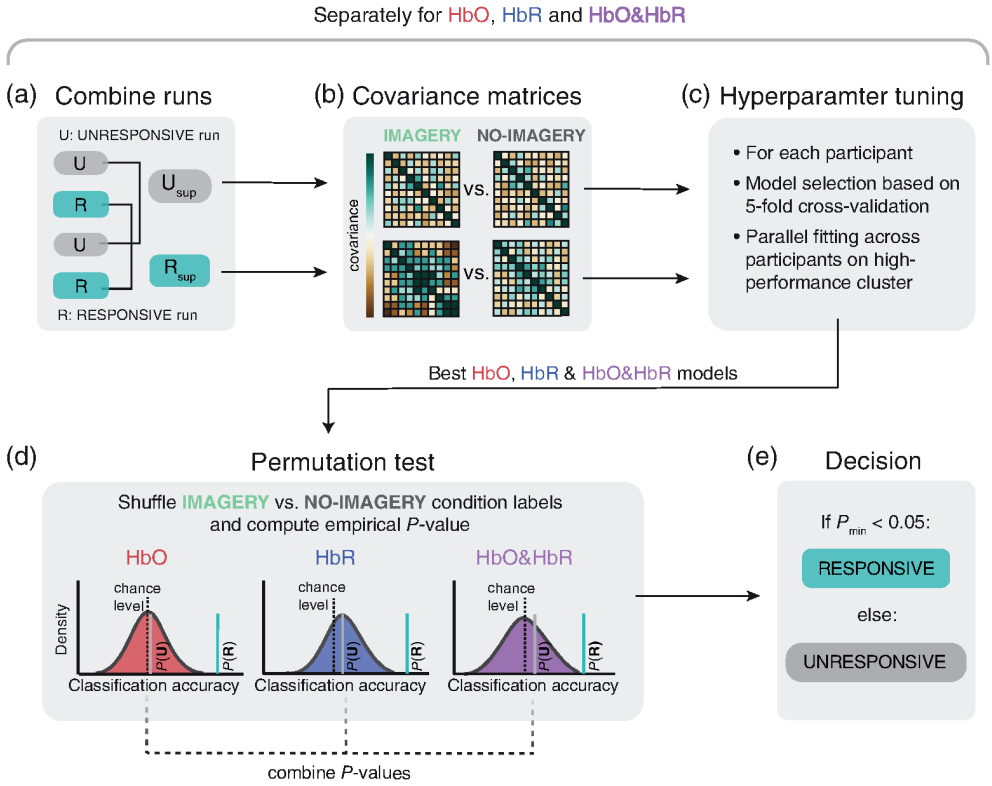

Figure 2. Covariance-based classification analysis.

A, The four individual runs were combined to "super runs" for subsequent analyses. **B**, Special-form covariance matrices were computed based on all channels and condition-specific trials per participant. **C**, For each participant, we conducted a small hyperparameter search to identify the best model, separately for HbO, HbR, and HbO&HbR. We determined the best model based on a 5-fold stratified cross-validation accuracy score. The procedure was conducted in parallel for all participants on a high-performance cluster. **D**. The best models' final classification accuracies were tested against the empirical null distributions with a permutation test (N = 250 permutations). **E**. In a final step, we classified the participant as RESPONSIVE if any P-value was below the threshold of 0.05, or else as UNRESPONSIVE.

Further statistical analyses

For all statistical tests, we set a significance threshold of $P < 0.05$, corrected for multiple comparisons. The statistical analyses described in this section were conducted in R Studio (Version 2021.09.1). Average classification accuracies are displayed as mean ± 1 standard deviation. Non-parametric Wilcoxon signed-rank tests were used to assess the difference between RESPONSIVE and UNRESPONSIVE for HbO&HbR, HbO, and HbR and task ratings. Spearman correlations were computed between, on the one hand, overall classification accuracies of RESPONSIVE/UNRESPONSIVE states and, on the other hand, either task ratings or performance ratings (i.e., alertness, concentration, and comfort). To assess the interaction effect of responsiveness (RESPONSIVE vs. UNRESPONSIVE) and time (RUN 1-4) on performance ratings, we computed a two-way ANOVA. Effect sizes for the *F*-statistics are reported by eta-squared.

Results

Covariance-based classification accuracy

When participants followed task instructions (during RESPONSIVE states), we were able to classify IMAGERY versus NO-IMAGERY trials above chance level. The average classification accuracy was highest for HbO&HbR (83.89% ± 6.93%), followed by HbR (82.22% ± 8.31%), and HbO (80.37%± 9.35%; Fig. 3C). When classification was based on HbO, accuracy exceeded the chance level for eight out of nine participants (all $P \leq 0.01$) and trended for participant P03 with an accuracy of 65% ($P = 0.08$). When HbR or HbO&HbR were used, all participants were classified correctly above chance level (all $P \leq 0.01$) in the RESPONSIVE state (Fig. 3A).

As expected, in the UNRESPONSIVE state, when the classifier did not receive two meaningfully distinct classes, average classification accuracy was around chance level for HbO (55.92 % ± 5.89%), HbR (55.56% ± 5.78%), as well as for HbO&HbR (55.93% ± 6.14%, Fig. 3C). Eight out of nine participants were classified at chance level for all three hemodynamic signals. Only one participant (P01) was inaccurately classified above chance level for HbO (66.67%, $P = 0.02$), HbR (68.33%, $P = 0.02$), and HbO&HbR (68.33%, $P = 0.01$, Fig. 3B).

We observed that the fNIRS-signal amplitudes at a single-trial level were generally higher in the RESPONSIVE than the UNRESPONSIVE state for both HbO and HbR in IMAGERY trials showing clear task-evoked brain responses only in the RESPONSIVE state (see example participant in Fig. S1). As expected, classification accuracy based on the covariance matrices of the fNIRS signals across channels was higher in the RESPONSIVE compared to UNRESPONSIVE state for HbO&HbR (Wilcoxon V = 45, $P < 0.01$), for HbO (V = 45, $P < 0.01$) and HbR (V = 45, $P < 0.01$; Fig. 3C). Similar results were obtained when averaging over HbO, HbR, and HbO&HbR for each participant (Fig. 3D). The average classification accuracy was 82.16% (± 8.40%) in the RESPONSIVE state, with a minimum of 71.67% (P03) and a maximum of 96.11% (P02). In contrast, the average classification

accuracy for the UNRESPONSIVE state was around chance level at 55.80% (± 5.51%). On average, classification accuracy in the RESPONSIVE state was higher than during the UNRESPONSIVE state (V = 43, $P < 0.01$, Fig. 3E).

To make a final judgment whether a participant was responsive or not, we took the minimum *P*-value obtained for HbO, HbR, and HbO&HbR. Based on this, we could correctly classify all participants as responsive, yielding 100% sensitivity (9/9 participants). For the UNRESPONSIVE state we obtained one false positive classification (P01), resulting in a specificity of 89% (8/9 participants). These results demonstrate that our covariance-based classification approach is highly sensitive (100%) and specific (89%) in detecting responsiveness to the instruction to perform a mental-imagery paradigm.

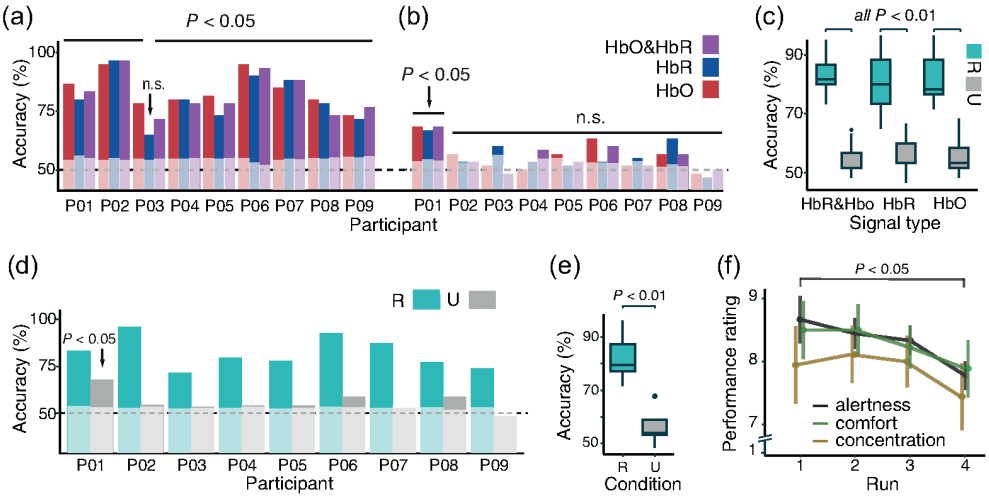

Figure 3. Classification accuracies and performance ratings.

A, Individual participants' classification accuracies for the RESPONSIVE state for HbO, HbR, and HbO&HbR separately. White shadings indicate individual empirical chance levels. All runs showed significant classification accuracies above the empirical chance level ($P < 0.05$), except the HbO data for P03. **B,** Same as **A** but for the UNRESPONSIVE state. For all participants, except P01 ($P < 0.05$), accuracies did not differ significantly from the empirical chance level. **C,** Mean accuracy over all participants, separately for HbO (R: 80.37%± 9.35% U: 55.92 % ± 5.89%), HbR (R: 82.22% ± 8.31%, U: 55.56% ± 5.78%), and HbO&HbR (R: 83.89% ± 6.93%, U: 55.93% ± 6.14%) for the RESPONSIVE and UNRESPONSIVE states (all $P < 0.01$)**.** During the RESPONSIVE state, classification accuracy was higher compared to the UNRESPONSIVE state for HbO, HbR, and HbO&HbR ($P < 0.01$). **D,** Average accuracy based on HbO, HbR, and HbO&HbR results. RESPONSIVE data were classified correctly based on $P_{min} < 0.05$. For the UNRESPONSIVE condition, no participant except P01 surpassed the chance level based on $P_{min} < 0.05$. **E,** Mean group accuracies combining HbO, HbR, and HbO&HbR (R: 82.16% ± 8.40%, N: 55.80% ± 5.51%; $P < 0.01$) **F,** Participants' mean alertness 8.47 ± 0.83, comfort 8.95 ± 0.97, and concentration 8.46 ± 0.92 ratings throughout the fNIRS runs. Alertness decreased significantly from run 1 to run 4 (Wilcoxon V = 30, $P < 0.05$).

User ratings and classification accuracy

To better assess the user experience of our paradigm, we acquired task ratings and performance ratings. Overall, the mental-imagery task was rated as easy (7.93 ± 1.24) and pleasant (8.81 ± 1,26) to perform. The mental-imagery task performance in the RESPONSIVE state was rated as being easier (Wilcoxon V = 31, $P = 0.04$) but not more pleasant (Wilcoxon V = 21.5, $P = 0.33$) than being passive in the UNRESPONSIVE state. Both, easiness and pleasantness of being RESPONSIVE or being UNRESPONSIVE were not associated with classification accuracies (all $P \geq 0.35$).

Furthermore, comfort, alertness, and concentration ratings during task performance were overall high (comfort: 8.95 ± 0.97, alertness: 8.47 ± 0.83, concentration: 8.46 ± 0.92) but followed a similar downward trend across runs **(**Fig. 3F). Specifically, alertness dropped significantly between runs 1 and 4 (Wilcoxon V = 30, $P < 0.05$). The small decrease in alertness was not correlated with a possible decrease in classification accuracy in the RESPONSIVE ($\rho = 0.28$, $P = 0.46$) or UNRESPONSIVE state ($\rho = 0.13$, $P = 0.74$). All three ratings were not significantly different between RESPONSIVE and UNRESPONSIVE across all runs (comfort: $F(3,24) = 0.64$, $P = 0.60$, $\eta^2 = 0.03$; alertness: $F(3,24) = 2.26$, $P = 0.10$, $\eta^2 = 0.02$; concentration: $F(3,24) = 0.62$, $P = 0.60$, $\eta^2 < 0.01$).

Discussion

This proof-of-principle study successfully implemented an fNIRS-based diagnostic tool for DoC patients in healthy adults. By employing command following through personalized mental imagery as a surrogate for MCS and UWS conditions, we demonstrated the potential of fNIRS for the differential diagnosis of DoC patients. Notably, the channel covariance-based classification achieved 100% sensitivity (9/9 participants) and 89% specificity (8/9 participants) in detecting participants' responsiveness, underscoring the potential of this highly sensitive and specific approach for clinical application.

A sensitive and specific diagnostic tool for bedside assessment

The development of a sensitive and specific bedside diagnostic tool is essential to advance the clinical assessment of DoC. In this study, we achieved correct classifications across almost all participants (except one). To the best of our knowledge, we are the first to show a 100% sensitivity in detecting volitional command-driven brain activity in healthy adults. In a recent study, Kazazian et al., (2024) reported a 62.5% sensitivity for a classical motor imagery task, and only 33.33% for spatial navigation. Our sensitivity is also higher than those reported in previous fMRI work (e.g. Edlow et al., (2021))—68.5%, Fernández-Espejo et al., (2014)—78.5%). In addition, we achieved this sensitivity using a paradigm that requires only 20 minutes (60 trials

per condition). This outcome suggests that our approach is highly effective with a relatively small trial number.

For bedside diagnostics, it is crucial to balance the number of trials needed for reliable classification with the cognitive demands placed on patients. To address this, we assessed participants' subjective ratings of concentration, comfort, and alertness, which remained high and stable across four experimental runs. While there was a slight but significant decrease in alertness from the first to the last run, this effect is likely attributable to the experimental setup. Our participants completed four runs to simulate both MCS and UWS conditions, which would not be necessary in clinical settings, where patients would be either responsive or unresponsive. Therefore, they would only need to complete half of the trials, reducing both task-switching effects and the overall test duration. Additionally, our testing environment — a small, dark, isolated cabin designed to minimize external noise and enhance signal quality — likely contributed to a decline in alertness over time.

Another central challenge in diagnosing DoC patients lies in the heterogeneity of individual brain-injury etiologies. Traditional analyses of hemodynamic responses often rely on a canonical HRF that assumes a uniform response shape. In contrast, our model-free covariance matrix approach does not depend on a predefined HRF, accommodating individual variability in hemodynamic response shapes. This flexibility is particularly valuable in the context of DoC, where patients exhibit diverse brain-injury patterns that affect brain response profiles (Irimia & Horn, 2015).

FNIRS measures two hemodynamic signals, namely HbO and HbR. We used HbO, HbR, and HbO&HbR channel covariances to separately classify datasets as RESPONSIVE or UNRESPONSIVE. In line with previous support-vector machine classification results in DoC (Li et al., 2021), we achieved the highest classification accuracy for the combined HbO&HbR data on a group level. This underscores the importance of leveraging both signals and their interactions to differentiate levels of mental engagement. Nevertheless, the combined HbO&HbR unequivocally yielded the maximum classification accuracy for only one individual participant, whereas either HbO (5/9) or HbR (3/9) alone yielded the best results for most participants on an individual level. Due to these individual differences in response characteristics, we based our responsiveness diagnosis on the combined contributions of HbO, HbR, and HbO&HbR, resulting in more robust predictions.

The one instance of misclassification (P01) in the UNRESPONSIVE condition raises important implications. A classifier trained on the data of this participant was able to significantly differentiate IMAGERY from NO-IMAGERY trials (HbO: $P = 0.02$; HbR: $P = 0.02$; HbO&HbR: $P = 0.01$), suggesting command following contrary to task instructions. This could indicate that the participant did not act unresponsively. Indeed, remaining at rest during the UNRESPONSIVE condition was experienced as more difficult by our participants than being active during the

RESPONSIVE condition. However, in contrast to our healthy controls, UWS patients who are truly unresponsive do not have to actively inhibit a response. Therefore, this misclassification should not arise in clinical diagnosis. Nevertheless, misclassification in general emphasizes potential limitations of relying solely on fNIRS data and underscores the value of integrating our paradigm into a broader diagnostic framework to minimize false diagnosis in DoC patients.

A novel and practical classification approach for fNIRS data

To the best of our knowledge, we are the first to show an fNIRS-based proof of principle in healthy participants of the use of Riemannian geometry for brain-based classification with direct clinical relevance. Covariance matrices are symmetric positive definite matrices and can be treated as points on a Riemannian manifold. We exploited this characteristic by performing classification in the tangent space of the channel covariance matrices using logistic regression. This approach can be computationally more efficient than conventional classification approaches, because each trial is represented by a single matrix rather than multi-channel HbO and HbR data that contain an additional temporal dimension. We further reduced the dimensionality of covariance matrices before classification by applying a spatial filter (Rivet et al., 2009). This approach demonstrates new opportunities to apply multivariate classification of fNIRS data in applied settings.

Given time and monetary constraints in clinical practice (Baal et al., 2018; Freedman et al., 2021), we aimed at making results readily available within less than 5 min after data acquisition while still maintaining high accuracy. This was specifically achieved by reducing covariance dimensionality, a small number of permutations (250), and a small parameter space during the hyperparameter search. The hyperparameter search was conducted to obtain a good fit for each classifier and was carried out for HbO, HbR & HbO&HbR, and each participant separately. The small parameter set reflected a trade-off between individually tailored classifiers and feasible fitting times (< 3 min). Noteworthy is that these durations may vary depending on the computational performance of the machine.

In most machine-learning applications, classifiers are trained on a separate training dataset and evaluated on a dedicated test dataset to assess generalizability. However, this split is unnecessary in our approach because our objective is not to evaluate classifier generalization but rather to use the successful training of classifiers as evidence of meaningful structure within the run. This structure, in turn, serves as an indicator for distinguishing between responsiveness and unresponsiveness during command-following tasks. To assess performance, we use cross-validation accuracy, which requires fewer trials than a traditional train-test split. This reduces the burden on patients with DoC and enables faster, clinically actionable results. Additionally, training individual models for each diagnostic session aligns with emerging trends in personalized medicine and diagnostics. In the context of DoC, using pre-trained models on large populations is likely not ideal due to the high heterogeneity of patients, symptoms, and impairments. For such

models to generalize effectively, they must account for substantial variability in fNIRS equipment (e.g., optode placements and preparation) and diverse clinical profiles. Our approach avoids this generalization challenge entirely. Here, classification is performed within specific sessions and does not need to generalize across different settings, patients. Instead, we demonstrate that Riemannian-based classification performs well with a small number of examples (Näher et al., 2025) and that retraining the model for each diagnostic run is both feasible and highly suitable for bedside clinical applications.

Benefits of a user-friendly mental-imagery task

Our paradigm is centered on a concise, user-friendly mental-imagery task, combining motor and speech imagery in 2 s trials. This design prioritizes minimal encoding time to reduce mental load and fatigue, potentially enhancing diagnostic accuracy in clinical applications. Previous studies using mental-imagery tasks employed significantly longer encoding durations, ranging from 12 s up to 30 s (Kempny et al., 2016; Kurz et al., 2018; Si et al., 2023). Despite the presumed advantage of these longer tasks for producing a robust hemodynamic response (HR), results have been mixed. Thus, our approach aimed to boost a robust HR with 2-s mental imagery. In a recent fNIRS-based BCI study, Vorreuther and colleagues (Vorreuther et al., 2023) successfully implemented this very short 2-s encoding time in a binary communication paradigm in healthy adults. Further research is required to confirm the suitability of this 2-s interval for DoC patients.

Mental-imagery tasks used in DoC diagnostics have varied in effectiveness, with motor-related tasks generally outperforming cognitively complex paradigms like mental arithmetic or spatial navigation (Kazazian et al., 2024; Kurz et al., 2018; Si et al., 2023). We implemented an individualized combined motor-speech imagery task which we think is tailored for DoC patients for two reasons:

Firstly, the task takes advantage of two modalities (i.e., motor and speech) that have both been shown to elicit robust HR in premotor/supplementary motor areas and Wernicke's area (Vorreuther et al., 2023). Since patients often have unique brain damage, capturing the activation of multiple cortical regions may increase the chance of observing task-evoked responses, even though some of the regions may not be fully functional anymore. Note that our optode layout was intentionally limited to the left hemisphere, primarily to ensure swift and practical placement which is essential for potential clinical use to reduce the burden on caregivers and patients. This left-lateralized montage may be optimal for capturing activation patterns related to speech and right-sided motor imagery – the dominant side for all our participants (see Table 1) – but not for left-handed individuals or those with atypical hemispheric dominance. Therefore, future studies might examine bilateral or hemisphere-tailored optode placements, potentially guided by handedness, to ensure optimal signal quality and generalizability.

Secondly, our mental-imagery task was individualized to the participants. Each participant was free to choose the specific strategy for the mental-imagery task they felt most comfortable with. In contrast to active paradigms currently used for DoC diagnosis (Abdalmalak et al., 2020; Kazazian et al., 2024; Kempny et al., 2016; Kurz et al., 2018; Molteni et al., 2013; Si et al., 2023), we did not give explicit task instructions but merely guided participants to choose a combined, motor-speech strategy with which they would try to get noticed if they were in the situation of an MCS patient. Giving this degree of freedom to patients allows adjusting the task to their mental capacities and may consequently increase the chance of obtaining a robust HR. The user-friendliness of the task is also reflected by consistently high ratings of easiness and pleasantness of our participants.

Limitations and future directions

The present study has some limitations that should be acknowledged. Overall, the sample size was relatively small (N = 9). While we acknowledge that this limits the generalizability of group-level findings and reduces statistical power to detect smaller or more nuanced effects, the primary aim of this study was not to establish population-level inferences but to demonstrate the feasibility of our individualized diagnostic paradigm on a single-participant level. Our approach relies on within-subject classification of mental imagery versus rest to assess responsiveness, which aligns with the needs of clinical diagnostics that require reliable detection at the level of individual patients. Indeed, each participant served as their own control, and classification accuracy was assessed via permutation-based significance testing at the individual level. Nevertheless, replication in a larger and more demographically diverse sample will be essential for future validation, particularly to evaluate classifier robustness across different physiological and cognitive profiles. A multi-center study or replication in a different lab with expanded recruitment would be a valuable next step toward clinical translation.

In addition to the sample size, our present sample has very limited diversity, including only one male participant and mostly young adults with lighter skin. To assess the generalizability of our paradigm across demographic variables (e.g. age, sex, handedness) and physical characteristics (e.g. hair density and type, and skin pigmentation) future studies should consider a stratified sampling approach. Given the relevance of physiological features for fNIRS signal quality (Wassenaar & Van den Brand, 2005), these variables could be quantified in a screening session (for guidelines see Yücel et al., 2024) and included in the analysis to assess their potential effects on the effectiveness of our classification approach.

One major challenge of our study design is the difficulty of enforcing unresponsiveness in healthy participants during simulated UWS conditions. In particular, the misclassification of one participant (P01) likely reflects a failure to fully comply with the instruction to ignore task cues, resulting in classifier detection of volitional mental activity. While UWS patients, by definition, are incapable of task engagement, healthy participants may inadvertently respond. This

limitation is inherent to a simulation-based paradigm. Although additional vigilance checks (e.g., eye tracking or monitoring of autonomic arousal) may provide indirect evidence of engagement, they cannot reliably detect covert mental activity. As an improvement to the current protocol, future studies could incorporate explicit post-run debriefings in which participants are asked whether they responded to any cues despite instructions, allowing researchers to exclude or repeat non-compliant runs. This would increase confidence in the validity of the simulated UNRESPONSIVE condition and reduce the risk of false-positive classifications.

A key limitation of the present study lies in the assumption that the responsive and unresponsive task conditions in healthy individuals can serve as proxies for the MCS and UWS diagnostic categories. This approach makes two strongly simplifying assumptions: first, that MCS and UWS are distinct and dichotomous clinical states, and second, that they can be meaningfully simulated in cognitively healthy individuals upon instruction. However, DoC is increasingly understood as a spectrum of awareness that cannot be fully replicated by healthy participants. Moreover, while participants were instructed to remain unresponsive, it is impossible to verify full suppression of task engagement, and this may differ from the neurological incapacity characteristic of UWS. Thus, our simulated conditions cannot fully reproduce MCS or UWS but rather simulate the presence or absence of volitional command following on an individual level. This study should therefore be understood as a proof-of-principle of the classification method under controlled conditions, which provides a foundation for future validation studies in actual DoC patient populations.

Future validation efforts to validate the proof-of-principle that we obtained in healthy adults for its intended use case should include testing our diagnostic tool with DoC patients. Concretely, we propose a validation study involving patients formally diagnosed with UWS and MCS using the current gold-standard behavioral assessments such as the Coma Recovery Scale-Revised (CRS-R). Patients would perform only the two RESPONSIVE runs of our individualized mental-imagery paradigm in its current form, focusing on whether the classifier can detect volitional imagery at the individual level. Patient outcomes should then be compared to behavioral assessments and, where available, EEG/fMRI indicators of covert awareness. The evaluation of classification sensitivity and specificity in a real-world clinical population, and the examination of feasibility in terms of recording time, patient tolerance, and processing speed are important next steps for translating our proof-of-principle into a diagnostic tool for bedside use.

Acknowledgements

This work was supported by The Netherlands Organization for Scientific Research (NWO; Vidi-Grant No. VI.Vidi.191.210 to BS) and by The Luik 3 grant for joint research between Cognitive Neuroscience and Knowledge Engineering Departments, Maastricht University, on advanced brain–robot interfaces (Grant No. 2015-9999) (to BS). GPT-4.5 was used for grammar checking.



4

Chapter 4

Primate Saccade Rhythmicity

This chapter is based on the following manuscript:

Näher, T.*, Zhang, Y.*, Spadacenta, S., Pandinelli, M., Bastian, L., Thier, P., Hafed, Z. M., & Fries, P. (2023). Primate saccade rhythmicity.

** These authors contributed equally.*

[Manuscript under review]

DOI: https://doi.org/10.1101/2023.09.27.559710

Abstract

Rhythmic sampling is a hallmark of many sensory systems in many diverse organisms. In primate vision, rapid foveating eye movements continuously scan the visual environment at a rate of 2-6 Hz, but the statistics of inter-saccadic interval distributions could equally well be described by non-rhythmic stochastic processes as by a (putatively variable) rhythmic oscillator. This raises the fundamental question whether primate visual scanning actually differs from the numerous examples of rhythmic sampling strategies exhibited across species and sensory modalities. Here, using experiments in humans, macaques, and a marmoset, as well as statistical approaches inspired by studies of temporal structure in neuronal spiking patterns, we show that primate saccade generation is unambiguously rooted in a rhythmic source. This finding was remarkably consistent across the three primate species despite their significant evolutionary distance. We find that saccade rate undergoes smooth, slow fluctuations. Accounting for those rate fluctuations was crucial for revealing the true degree of saccade rhythmicity. Thus, exact saccade timing is determined by an interaction between the slow saccade-rate fluctuations with the more temporally-local rhythmic generator. This demonstration of rhythmicity in overt oculomotor sampling behavior provides a link to the fundamental rhythmic nature of many central perceptual and cognitive processes, and places primate saccadic sampling into the large family of rhythmic active sampling behaviors.

Introduction

Vision is an active process: rather than passively receiving stimuli, organisms actively reposition their eyes to selectively sample behaviorally relevant information from the environment. In primates, saccadic eye movements shift the fovea to successive locations, creating a sequence of high-acuity "snapshots". During natural viewing, saccades occur repetitively, with inter-saccadic intervals (ISIs) of a few hundred milliseconds, yielding mean saccade rates of 2–6 Hz (Otero-Millan, Macknik, et al., 2013). This repetitive sampling bears a striking resemblance to other active-sensing behaviors such as whisking and sniffing in rodents (Kleinfeld et al., 2006). Those behaviors are known to be strongly rhythmic, raising the question of whether similar intrinsic rhythmic processes also govern saccade generation (Takatoh et al., 2022).

Several influential accounts of saccade behavior have, in fact, implicitly contained or predicted some degree of rhythmicity in the saccadic system. Classical descriptions of the temporal structure of saccades emphasized two temporal regularities: a refractory period following each saccade, during which the probability of initiating the next movement is suppressed, and a unimodal shape of the ISI distribution (Noorani & Carpenter, 2016). Both of these saccade characteristics can be explained by a saccade generation model based on a repetitive rise-to-threshold accumulation. Together, this implies that saccades show distinct temporal structure.

However, detailed inspection of saccade timing patterns reveals large variability in observed ISIs across many behavioral tasks, resulting in heavily skewed ISI distributions with substantial tails produced by fairly long ISI values. Such skew significantly complicates the interpretation of saccade timing as a rhythmic process, with some accounts even going as far as arguing that skewed ISI distributions are a manifestation of a non-rhythmic stochastic process governing saccade generation (Amit et al., 2017). This, if true, would challenge prevailing assumptions about the strong links between eye movements and cognitive processes like attention (Fiebelkorn et al., 2013; Fries, 2023; Landau & Fries, 2012; Tian et al., 2016) and visual sensitivity (J. Bellet et al., 2017; Benedetto & Morrone, 2017), which are themselves believed to be strongly rhythmic. Therefore, demonstrating, and understanding the properties of saccade-system rhythmicity is a fundamental problem in primate systems neuroscience.

In this work, inspired by statistical analysis approaches historically used to study temporal dynamics in neuronal spiking, we demonstrate genuine saccade-system rhythmicity in the primate brain. Using experiments with humans, macaques and a marmoset performing the same visual tasks, we show that high rhythmicity can still give rise to skewed ISI distributions when coupled with slow fluctuations in saccade rate, putatively due to factors like arousal, scene statistics, or task context. Critically, differences in apparent rhythmicity across tasks and species are driven by how neuronal circuits adjust saccade rate in response to task demands, not by changes in the regularity of the underlying generator. These findings reconcile previous

conflicting accounts of saccade timing and demonstrate that a locally rhythmic generator does exist for saccades. Beyond providing a link between rhythmicity in covert attentional sampling and overt saccadic sampling, our results motivate investigation into the neuronal mechanisms underlying rhythmic saccade drive.

Methods

Participants and animals

Two adult rhesus macaques (*Macaca mulatta*, denoted C and G in the main text, both 19 years old) participated in this study. Both animals were implanted with a titanium headpost. Monkey G had a recording chamber implanted for addressing other scientific questions. To maintain motivation during the behavioral task, a small juice reward was given after each correct trial.

One adult marmoset (*Callithrix jacchus,* 6 years old) participated in the study.

Twenty-five adult human participants (12 males, 13 females), recruited from the public, participated in the current study. We excluded three participants (all female) due to insufficient number of trials. All human participants had normal vision and provided informed consent before participation.

Tasks

Figure 1 illustrates the stimuli. Importantly, all species performed the same tasks with only minor differences in the reward procedure (auditory feedback for humans versus juice delivery for macaques and marmoset).

Macaque task

Macaques were recorded in two types of sessions. Each type of session contained a specific set of tasks.

Session type 1 consisted of the REGULAR and IRREGULAR tasks. Each trial in both tasks began with the animal fixating at the center of the screen on a fixation dot (diameter 0.1 dva, i.e. degrees of visual angle). After fixation for 800 to 1000 ms, the fixation dot was turned off, followed by a ≈100 ms gap. The search array was then presented. The array consisted of white filled dots (diameter 0.3 dva) that reflected potential saccade targets (see Figure 1 for an illustration). Within the search array, 20% of the presented dots were real targets. The animal's task was to sequentially look at the dots to find any one of the real targets (foraging). During the foraging, if the animal's gaze landed inside a fixation window around a real target (radius 1.0 dva, not visible to the monkey), that target was replaced by an image of grapes, and a small juice reward was provided. If the monkey kept fixation on the grapes, multiple pulses of reward were

provided every 200 ms until a maximum window of 2000 ms had passed. This was followed by an inter-trial interval of 500 ms and the start of a new trial with the presentation of a new fixation point. If the maximum search time passed (5000 ms) and none of the real targets had been foveated, the trial was terminated without a reward, followed by the same inter-trial interval and the start of a new trial. In the REGULAR task, the dots in the search array were spaced as if they were the centers of an invisible hexagonal tiling. In the IRREGULAR task, the search array contained the same number of dots as in a REGULAR trial, but the positions of the dots were randomly redistributed across the full screen to break any spatial regularity.

Session type 2 consisted of only one task, the IMAGE task, and began with the animal fixating at the central fixation dot. This pre-search fixation for the IMAGE task was identical to the REGULAR or IRREGULAR tasks. After fixation for 800 to 1000 ms, the fixation dot was turned off, and after a ~100 ms gap, it was replaced by a natural image in color (16 * 21 dva, McGill Calibrated Colour Image Database, http://tabby.vision.mcgill.ca). From the moment the fixation dot was turned off, the animal was free to explore the image. After a random delay of 100 ms to 1900 ms, a small white dot (either 0.1 dva radius or 1.7 dva radius, randomly interleaved), the target, appeared at a random location within the image. We used two different radii for a comparison that is not relevant to the current study. The monkey's task was to locate this target within a period of 5000 ms. If the target dot was found, the reward procedure as described for session type 1 followed.

Marmoset task

The marmoset task closely mirrored the macaque version, with parameter adjustments to accommodate species-specific behavioral and attentional constraints. The key differences are described below.

All three tasks (REGULAR, IRREGULAR, and IMAGE) were presented within a single session type, in random order. Each trial began with fixation on a central red dot (diameter 0.4 dva) which the animal had to acquire within 200 ms and hold for up to 600 ms. The fixation dot was then turned off and followed by a ~100 ms gap. In the REGULAR and IRREGULAR task, the search array was then presented. Array composition, target proportion (20%), dot spacing rules, and maximum search time (5000 ms) were identical to the macaque version.

The fixation window around real targets was enlarged (2 x 2 dva, not visible to the monkey) to reduce frustration and encourage longer working times. Upon successful target fixation, the target was replaced by an image of a marmoset face (3 × 4 dva), and a small amount of marshmallow juice was delivered. If the monkey kept fixation on the marmoset face, short pulses of reward were provided manually until trial end.

For the IMAGE task, stimuli (16 x 21 dva) were drawn from the same database. Only a single target size was used (diameter 0.3 dva, white). The monkey's task was to locate this target within a period of 5000 ms. If the target dot was found, multiple pulses of reward were manually provided until a maximum window of 3000 ms had passed.

Note that due to a bug in the stimulation script, the end-of-trial marker was missing in the marmoset data. We therefore decided on a fixed trial length of 4 s in which most marmosets found the reward.

Human tasks

Human participants were recorded in two types of sessions. Each type of session contained a specific set of tasks.

Session type 1 consisted of the REGULAR, IRREGULAR, and IMAGE tasks. These tasks for human participants were identical to the respective tasks for macaques, except for the following differences: (1) The pre-search fixation duration was 2500 to 3500 ms; (2) the post-search fixation duration was also 2500 to 3500 ms and did not contain a repetitive reward; (3) in all tasks, successful trial completion was indicated by a short audio feedback. If the maximum search time passed and no target was found, the trial was terminated. Inter-trial-intervals were 2000 ms.

Session type 2 included the REGULAR and IMAGE tasks, as well as the free-viewing MOVIE task. It was identical to session type 1, except for the following differences: (1) The REGULAR task in session type 2 contained three different spacings of the grid (hexagon radius of 2.3, 4.5, and 8.5 dva); (2) the IMAGE task remained unchanged from session type 1 despite a different set of images being used.

The free-viewing MOVIE task (performed exclusively by human participants) consisted of a 10-minute naturalistic first-person-view video. The video contained minimal cuts, such that uninterrupted epochs lasted several minutes. Furthermore, the camera trajectory was smooth, minimizing transients due to camera movement. The movie is in color, was downloaded from YouTube (https://www.youtube.com/watch?v=hld4uaO1MDE) and can be obtained from the authors upon request. We chose this movie, because it provides naturalistic visual content, motivates subjects to sustained saccadic exploration of the continually renewing input, and largely avoids transients. Participants were instructed to keep their head still on the chinrest and simply watch the movie without performing a specific task.

Eye movement recording

All visual stimulation was controlled by custom software (https://github.com/esi-neuroscience/ARCADE) and presented on an LCD monitor. For macaques, the monitor was an LG 32GK850G-B

with a refresh rate of 143.9 Hz. For marmoset monkeys, the monitor was a 10 Inch Beetronics with a refresh rate of 60 Hz. For human participants, the monitor was a ViewPixx (VPixx Technologies Inc) with a refresh rate of 120 Hz. In humans and macaques, we recorded binocular gaze data at 500 Hz using an EyeLink-1000 system (SR Research Ltd.). For marmoset monkey, we recorded monocular gaze data at 120 Hz using an ISCAN ETL-200 infrared eye tracking system. All gaze data were up-sampled to 1 kHz to achieve a common sampling rate for all further analyses. Real-time gaze data from one eye was utilized for online behavioral control such as the pre- and post-search fixation periods.

The human data were collected across one or two sessions, with some participants attending only the first session and others attending both. The first session included a brief training block to ensure participants gained experience with the online eye-tracking control and to fully understand the task. A given session encompassed up to six blocks with the REGULAR, IRREGULAR, and IMAGE tasks randomly interleaved. Each block consisted of 60 trials, 20 of each task. Participants worked on one block until 60 correct trials (20 of each condition) were achieved. Participants were encouraged to complete all six blocks; however, the total number completed was contingent upon each participant's level of fatigue. The experimental design sought to prioritize participant comfort and data accuracy over the volume of completed blocks. The second session used the same approach, but for session type 2.

Saccade detection

We used a convolutional neural network (CNN) for saccade detection (M. E. Bellet et al., 2019). We trained separate networks for humans, macaques, and the marmoset on randomly selected trials from each species. First, we labeled saccades and fixations as separate classes in the data using a custom GUI. After epoching the data into trials, we replaced blinks with NaNs. Subsequently, we used those preprocessed data to train the CNN, then detected saccades for each eye separately. We only accepted saccades if they occurred in both eyes simultaneously. For all detected saccades, we enforced a minimum duration of 6 ms and a minimum separation between two subsequent saccades of 15 ms.

Trial-wise saccade rates

For all further analyses of the trial-based conditions (i.e. REGULAR, IRREGULAR, IMAGE, but not MOVIE), only trials with a minimum duration of 1 s were included.

A central aspect of the analysis will be concerned with saccade rates. For the trial-based conditions, we computed the mean saccade rate for each trial and used this to quantify between-trial variability in saccade rate. Saccade rate was computed as the total number of detected saccades within the trial divided by the trial duration in seconds, after excluding the initial 100 ms post stimulus onset, because saccade rates are known to show a transient decrease after fixation or stimulus onset (Reingold & Stampe, 2000; Engbert & Kliegl, 2003;

Fries et al., 2008; Buonocore & Hafed, 2023). Trial-wise saccade rates were used for all analyses relating saccade rate to the rate-corrected rhythmicity index (RCRI, see below), including rate stratification and regression analyses.

Gamma renewal modeling with time-varying rate

To investigate whether heavy-tailed inter-saccadic-interval (ISI) distributions can arise from a locally regular (i.e. rhythmic) generator operating under time-varying saccade rate, we fitted a gamma renewal model in which the shape of the ISI distribution is fixed within a subject, while the mean of each ISI distribution is determined by the trial-wise rate. The details of such fitting are provided below.

For each subject (humans, macaques, and marmoset), we extracted all valid ISIs. For model fitting and visualization, all analyses used ISIs within 100–600 ms range, thereby matching the range that we used for histogram-based comparisons.

Local rate estimation

To capture the rate variability, we estimated the local saccade rate associated with each ISI using a sliding window of width $W = 1$s. Let t_i and t_{i+1} denote the onset times (in seconds) of two successive saccades, and $\Delta t_i = t_{i+1} - t_i$ the corresponding ISI. We assigned to each ISI a midpoint time $m_i \ = (t_i + t_{i+1})/2$. The local rate r_i was computed as the number of saccades whose onset times fell within $[m_i - W/2,\ m_i + W/2]$ divided by the actual window duration after truncation at the trial boundaries.

Gamma renewal model with rate-dependent mean

We modelled each Δt_i, i.e. each ISI, as an independent draw from a gamma distribution with a single subject-specific shape parameter k and a rate parameter β_i determined by the local rate r_i. Specifically, the model imposes a local mean interval

$$\mathbb{E}[\Delta t_i] = \frac{1}{r_i}$$

and uses the gamma parameterization in which mean $= k/\beta_i$. Thus, the rate parameter varies across ISIs as

$$\beta_i = k r_i$$

To estimate k without making assumptions about the precise time course of the local rate, we exploited the time-rescaling theorem. If the model is correct, the rescaled variables

$$z_i = r_i \cdot \Delta t_i \sim \mathrm{Gamma}\,(k, k)$$

regardless of how r_i varies across ISIs. We therefore estimated the local rate at the trial level as $\hat{r} = n_{\text{filt}}/T_{\text{filt}}$, where n_{filt} is the number of ISIs and T_{filt} is their total duration in seconds. We then formed $z_i = \hat{r} \cdot \Delta t_i$ for all ISIs within that range and fitted a gamma distribution to the pooled z values by maximum likelihood yielding $\hat{k}$ and its standard error from the Fisher information. Goodness of fit was assessed with percentile–percentile (PP) plots of the empirical z values against the fitted $\text{Gamma}\left(\hat{k}, \hat{k}\right)$ distribution, as well as formally with a one-sample Kolmogorov–Smirnov (KS) test comparing the empirical distribution of z to the fitted gamma cumulative distribution function (CDF); a non-significant KS statistic indicates that the data are consistent with the model.

To interpret the fitted shape parameter, we define "high" k in terms of the skewness of the local gamma ISI distribution. For a $\text{Gamma}\,(k, \beta_i)$ distribution, the skewness depends only on the shape and is given by

$$\text{Skew}\,(\Delta t) = \frac{2}{\sqrt{k}}.$$

We take "mildly skewed" to mean $\text{Skew}\,(\Delta t) \leq 0.75$, which implies

$$\frac{2}{\sqrt{k}} \leq 0.75 \quad \Rightarrow \quad \sqrt{k} \geq \frac{8}{3} \quad \Rightarrow \quad k \geq \frac{64}{9} \approx 7.1.$$

Thus, throughout we refer to $k \geq 8$ as a high gamma shape parameter under this criterion, corresponding to a locally only mildly right-skewed ISI distribution.

Rate-corrected ISI distribution

To visualize the structure of ISI variability after accounting for trial-level rate fluctuations, we computed a rate-corrected ISI for every empirical ISI. For each $\Delta t_{i,j}$ (the *j*-th ISI in trial *i*) with associated local rate estimate r_i, we formed the dimensionless ratio

$$z_{i,j} = r_i \cdot \Delta t_{i,j}$$

which expresses each ISI as a multiple of the expected interval $1/r_i$ at that local rate. Under the gamma renewal model, $z_i \sim \text{Gamma}\,(k, k)$ with mean 1, so this rescaling removes between-trial rate differences while preserving within-trial ISI variability. Rescaled ISIs were pooled within each species and their distributions visualized as kernel density estimates.

Rate-corrected rhythmicity measure

To quantify local rhythmic regularity while minimizing sensitivity to slow rate fluctuations, we used a rate-corrected rhythmicity index, $RCRI$, based on the local variation of successive ISIs.

Local variability was quantified using the CV_2 statistic, defined for successive ISIs as

$$CV_2\,(i,j) = \frac{2\ \left|\Delta t_{i,j+1} - \Delta t_{i,j}\right|}{\Delta t_{i,j+1} + \Delta t_{i,j}}.$$

The overall local variation was computed as the mean across intervals, $CV_2 = \langle CV_2(i) \rangle$. The rate-corrected rhythmicity index was then defined as

$$RCRI = \frac{1}{1 + CV_2}.$$

For a Poisson process with exponentially distributed ISIs, the expected value of CV_2 is 1, yielding $RCRI = 0.5$. Values of $RCRI > 0.5$ therefore indicate greater local regularity than expected from a Poisson process. Furthermore, RCRI is analytically scale-invariant: because CV_2 is a ratio of differences to sums of consecutive ISIs, multiplying all intervals by a constant c leaves the index unchanged (the factor c cancels in numerator and denominator), making it independent of absolute saccade rate.

The CV_2 statistic and related local variation measures are widely used in neuronal spike-train analysis to quantify firing regularity under non-stationary conditions, where mean firing rates may vary over time (Holt et al., 1996). Because CV_2 depends only on pairs of successive intervals, it is less sensitive to slow changes in mean rate.

Stationary vs non-stationary processes and spectral measures

To isolate the effects of slow rate fluctuations on rhythmicity metrics and spectral measures, we simulated event sequences using stationary and non-stationary gamma renewal processes that were matched in local regularity but differed in their rate dynamics.

In both simulations, saccade times were generated as a gamma renewal process, in which successive ISIs were drawn independently from a gamma distribution. Local regularity was held constant by fixing the coefficient of variation of the ISI distribution to $CV = 0.15$, corresponding to a highly regular, oscillator-like generator.

In the stationary process, the mean ISI was fixed by setting a constant rate equal to the mean of an empirical rate time course derived from the MOVIE task, resulting in a time-invariant ISI distribution.

In the non-stationary process, the mean ISI varied over time according to that same empirical rate time course from the MOVIE task. At each step, the rate of the gamma process was set to the inverse of the current empirical rate, while the shape of the distribution, and thus local regularity, was kept constant. As a result, ISIs were locally regular but globally heterogeneous due to slow rate fluctuations present in the empirical rate dynamics.

Spectral analysis of event trains

To quantify rhythmicity with spectral measures, simulated event times were discretized into a binary event train at 1 ms resolution (sampling frequency 1000 Hz). Power spectral density (PSD) estimates were computed using Welch's method (segment length 2048 samples, constant detrending, density scaling).

For each simulation, we extracted the maximum PSD value within the 2–6 Hz frequency range ("peak power") as a summary measure of spectral rhythmicity (Amit et al., 2017).

We generated multiple (n = 500) independent realizations of the stationary and non-stationary processes. For each realization, we extracted the peak power in the 2–6 Hz range and the $RCRI$ computed from the simulated ISIs.

Rate-stratified analysis of saccade rhythmicity

To assess the relationship between saccade rate and rhythmic structure using a nonparametric approach, we performed a rate-stratified analysis within each species separately. For each trial, mean saccade rate was computed as defined earlier.

Within each species, trials were ranked by mean saccade rate and assigned to tertiles (low, medium, high) using quantile binning, ensuring equal numbers of trials per bin. This procedure controlled for differences in saccade rate across species and allowed rate-dependent effects to be examined within comparable relative rate ranges.

For each species and rate bin, saccade autocorrelation functions were computed and then averaged across trials within each bin. Additionally, ISI distributions were obtained by pooling ISIs across trials within the bin. Within each species, trial-wise rate-corrected rhythmicity values (as defined above) were grouped by rate bin.

For humans, pairwise comparisons of rhythmicity between rate bins were performed using independent-samples *t*-tests. For macaques and marmoset, where trials were pooled across individuals (Fries & Maris, 2022), pairwise comparisons were performed using two-sided Mann–Whitney U tests. All statistical tests were conducted within species, and significance was assessed at $\alpha = 0.05$.

Within-trial slow rate fluctuations and induced higher-order dependencies

Windowed saccade-rate estimates

Within each MOVIE trial, saccade onset times were binned into non-overlapping windows of 3 s. For each window, saccade rate was computed as the number of saccades divided by window duration (Hz), yielding a within-trial rate time series.

Blinks were identified as time segments in which eye-position samples contained NaNs. For analyses based on ISIs, any ISI which overlapped a blink segment was excluded.

For each trial, we computed the lagged serial correlation of the windowed saccade-rate time series for lags up to 20. Serial correlation functions were then averaged across trials.

To characterize the frequency content of within-trial rate fluctuations, we estimated the power spectral density of the windowed rate series using Welch's method. Because rate was sampled once per 3 s window, the effective sampling frequency was 1/3 Hz.

Monte Carlo null model for rate autocorrelation

To test rate autocorrelations against a memoryless baseline, we constructed a per-trial surrogate under a homogeneous Poisson null while preserving each trial's total saccade count. For each eligible trial with N_i saccades and B_i windows, surrogate window counts were generated via a multinomial draw with uniform bin probability $1/B_i$. Surrogate rates were computed as counts divided by 3 s. The autocorrelations were computed identically to the empirical data.

This procedure was repeated 500 times. For each lag, we derived a 95% null envelope based on the 2.5th–97.5th percentiles.

ISI serial correlations

For each trial, we again computed the ISI sequence as the difference between successive saccade onset times (see above). Higher-order temporal dependencies were quantified using lagged serial correlations across the ISI sequence: for lag k, the correlation between $ISI_n$and ISI_{n+k} was computed for $k = 1, \dots, 20$. Correlation functions were averaged across trials.

Rate-conditioning via integrated intensity transformation

To assess ISI serial correlations after conditioning on within-trial rate fluctuations, we applied a time-rescaling transformation using the trial's windowed rate estimate. This rescaling allows us to analyze ISI time courses, while removing correlations introduced by rate fluctuations.

1. The windowed rate series was linearly interpolated to obtain a continuous rate function $\lambda(t)$, with $\lambda(t) \geq 0$.

2. For each ISI spanning $[t_n, t_{n+1}]$, we computed the integrated intensity

$$\Lambda_n = \int_{t_n}^{t_{n+1}} \lambda(t)\, dt,$$

using numerical integration with a step size of 0.01 s. Importantly, Λ_n gives the duration of the interval $[t_n, t_{n+1}]$ measured in rescaled (intensity) time. This transformation expresses each

ISI in units of expected event count, thereby removing correlations induced solely by slow rate fluctuations.

3. Serial correlations were then recomputed on the transformed interval sequence $\{\Lambda_n\}$ using the same lag range (1–20).

Drift-to-threshold model

To formalize the relationship between saccade rate and rhythmicity, we implemented a drift-to-threshold model (Figure 7; similar to LATER models)(Noorani & Carpenter, 2016). In this model, a decision variable accumulates linearly with drift rate μ and additive Gaussian noise (diffusion coefficient σ) toward a fixed threshold ϑ. Upon threshold crossing, a saccade is triggered, and the accumulator resets to zero, followed by a fixed refractory period τ. The resulting inter-saccade intervals follow an inverse Gaussian distribution with mean ϑ/μ and shape ϑ^2/σ^2, shifted by τ. We fixed the threshold at $\vartheta = 1$ (arbitrary units) and the refractory period at $\tau = 50$ ms and fitted the single remaining free parameter σ to the scatter plot of RCRIs as function of empirical rates, as presented in Figure 4, by minimizing the sum of squared errors. For each empirical saccade rate r, the corresponding drift was determined analytically as $\mu = \vartheta / (1/r - \tau)$. Predicted RCRI values were computed deterministically via numerical integration over the inverse Gaussian probability density function on a 500-point grid. To illustrate how natural rate fluctuations affect rhythmicity, we additionally simulated an accumulator driven by a time-varying drift signal with $1/f$ spectrum, producing a continuous mixture of saccade rates between 1.5 and 4.8 Hz.

Statistical analyses

To assess whether trial-wise saccadic rhythmicity could be explained by saccade rate, we performed species-specific linear regression analyses using ordinary least squares. The dependent variable was trial-wise RCRI, as defined above. Mean saccade rate (Hz) was computed for each trial as the number of saccades divided by trial duration.

For each species, rhythmicity was modelled as a function of mean saccade rate and experimental task. Task was included as a categorical predictor with three levels (REGULAR, IRREGULAR, IMAGE).

For humans, additional comparisons of rhythmicity across rate-defined bins were performed using independent-samples *t*-tests. For macaques and marmosets, where data were pooled across individuals, comparisons were performed using two-sided Mann–Whitney U tests.

Unless otherwise stated, all statistical significances were assessed at $\alpha = 0.05$. All analyses were conducted in Python 3.11 using NumPy (v1.24), pandas (v2.0), SciPy (v1.11), and *statsmodels* (v0.14).

Ethics statement

Humans

The study was approved by the local ethics committee (Ethikkommission des Fachbereichs Medizin der Goethe-Universität, Nr 2022-1031_1-). Participants gave written consent before participating in the study.

Macaques

All procedures and housing conditions complied with the German and European law for the protection of animals (EU Directive 2010/63/EU for animal experiments). All surgical and experimental methods were approved by the regional authority (Regierungspräsidium Darmstadt) under the following permit numbers: F149-1008 & F149-1007. To support their wellbeing, animals were offered dietary variety, habitat enhancements, and interactive objects.

Marmosets

All experimental protocols were approved by the local animal care committee (Regierungspräsidium Tübingen, Abteilung Tierschutz) and fully complied with the German and European law for the protection of animals (EU Directive 2010/63/EU for animal experiments).

Results

We characterized the nature of saccade rhythmicity in three primate species by having them view static or dynamic scenes (Figure 1). In each case, we assessed to what extent the timing between successive saccades (ISI) could be thought of as being generated by a rhythmic process, and we did so by relating such timing to slower saccade rate fluctuations. We hypothesized that ISI distributions can be effectively modeled as a doubly-stochastic process, with one component governed largely by a relatively regular oscillator, and another component that slowly alters the period of the oscillator (and thus the overall saccade rate). In what follows, we first show that saccade rate substantially fluctuates in all three species, and we then explain how accounting for such fluctuations reveals the degree of intrinsic rhythmicity in saccades.

Saccade rates vary across trials

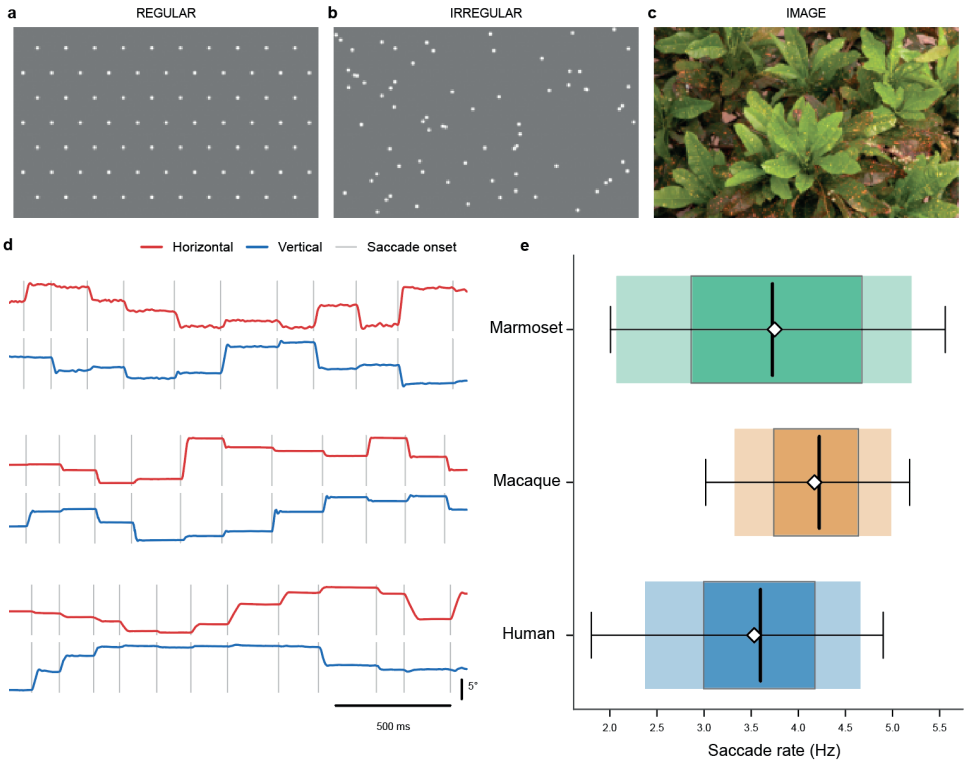

Figure 1. Saccade rates vary across trials.
(a-c) Example task screens of the three conditions: REGULAR, IRREGULAR, and IMAGE. These conditions were used for all analyses in this figure. **(d)** Example eye traces from all three species. Detected saccade onsets are indicated by vertical lines. Temporal scale bar indicates 500ms, vertical scale bar 5 degree of visual angle. **(e)** Box plots reflecting the distributions of saccade rates over trials. Trials were pooled over conditions (REGULAR, IRREGULAR, and IMAGE) and subjects, separately for the three species. Median [interquartile] rates were 3.60 [3.00, 4.18] Hz (human, n = 22 subjects), 4.22 [3.74, 4.64] Hz (macaque, n = 2 subjects), and 3.73 [2.86, 4.68] Hz (marmoset, n = 1 subject). The 5th–95th percentile ranges were 1.80–4.90 Hz, 3.02–5.18 Hz, and 2.01–5.56 Hz, respectively. The darker shaded box spans the interquartile range (25th–75th percentile), the lighter shaded region extends from the 10th to 90th percentile, the vertical black line marks the median, the white diamond indicates the mean, and the whiskers with caps denote the 5th–95th percentile range. Mean ± SD of the saccade rates were 3.53 ± 0.95 Hz in humans, 4.17 ± 0.67 Hz in macaques, and 3.75 ± 1.18 Hz in marmosets. Saccade rate was computed per trial as the number of saccades divided by trial duration (s), after excluding the first 100 ms following array/image onset to minimize onset-transient effects.

We first quantified trial-to-trial variability in saccade rate across trials for the three investigated species: human, macaque and marmoset. For each species, saccade rates varied widely rather than clustering tightly around a single characteristic value (Figure 1e). In humans, trial-wise

rates spanned nearly 3 Hz across the central 90% of trials (mean ± SD: 3.53 ± 0.95 Hz, 5th–95th percentile: 1.80–4.90), indicating substantial variability. Macaques showed a similarly broad distribution (mean ± SD: 4.17 ± 0.67 Hz, 5th–95th percentile: 3.02–5.18), while the marmoset exhibited the largest relative spread (mean ± SD: 3.75 ± 1.18 Hz, 5th–95th percentile: 2.01–5.56). Thus, in all three species, saccade rates varied substantially from trial to trial, indicating that the underlying process cannot be modeled with a single stationary rate.

Gamma renewal models can reproduce heavy-tailed ISI distributions

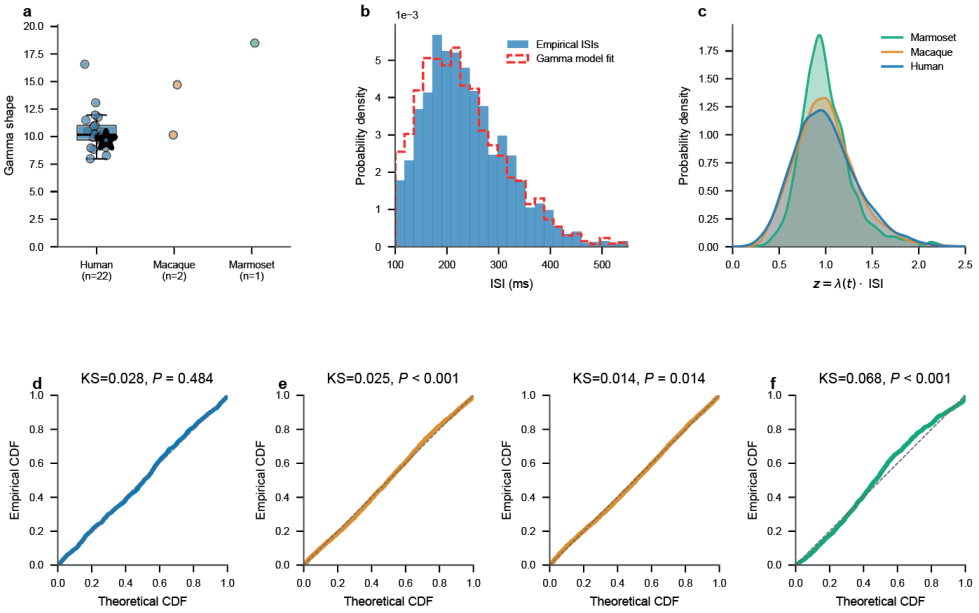

Figure 2. Gamma renewal models can reproduce heavy-tailed ISI distributions.
Saccade rates were determined empirically per trial and used to normalize the trial-wise ISIs. The normalized ISIs were pooled over trials and conditions, per participant, and the resulting distributions were fitted with a gamma renewal model with the shape parameter as the only free parameter. **(a)** Fitted gamma-shape parameters. Points indicate individual participants; the star indicates the example participant in panel b. The overall large shape-parameter values indicate strong local regularity: after accounting for trial-wise saccade rate, the rate-normalized ISIs are near-symmetric (skew ≤ 0.75). **(b)** Blue: Empirical ISI distribution from the example human participant. Red: Best-fitting gamma-renewal model; The model reproduces the heavy-tail ISI distribution, despite being locally regular with a gamma shape parameter value close to ten. **(c)** Rate-corrected residual kernel density ISI distributions after correcting per trial for the expected value of ISI, separately for all three species. **(d)** Example human probability-probability (PP) plot, showing that rate-corrected ISI distributions follow the proposed gamma distribution shape. For humans, all KS-tests were non-significant, indicating that the gamma model is a good fit. **(e)** same as d, but for the two macaques. **(f)** same as in e, but for the marmoset. Note that the KS tests were significant for the macaques and the marmoset.

The large trial-to-trial variability in saccade rate observed in all tested species (Figure 1) has two important implications for the interpretation of saccade timing statistics. First, the empirically observed heavy-tailed inter-saccadic-interval (ISI) distribution reflects a marginal distribution aggregated over time, rather than the interval statistics generated at any fixed moment. When saccade rate varies, the overall ISI distribution becomes a mixture of many locally generated interval distributions with different means, such that heavy tails can arise even if saccades are generated by a locally regular process. Second, rate fluctuations can introduce higher-order temporal structure that can affect spectral analyses of saccade rhythmicity.

To test these implications explicitly, we modeled saccade generation as a doubly stochastic gamma renewal process, with temporally local stochasticity and a time-varying rate. These saccade rates were determined empirically per trial. Using the time-rescaling theorem, we rescaled each ISI by its trial-wise rate to obtain normalized ISIs. The normalized ISIs were pooled over trials and conditions, per participant. The resulting distributions were fitted with gamma distributions, where the only free parameter was the shape parameter k.

In humans, fitted k values were 10.49 ± 1.78 (5th–95th percentile: 8.31–13.02; min–max: 7.98–16.57). Macaques showed similar values (12.43 ± 2.28; 5th–95th percentile: 10.38–14.48; min–max: 10.15–14.71). The marmoset data yielded a gamma shape parameter of 18.49. All fitted values exceeded the threshold of $k \geq 8$, corresponding to a local ISI distribution with skewness ≤ 0.75, indicating that distributions were only mildly right-skewed. Goodness of fit was assessed visually with probability-probability (PP) plots of the rescaled ISI distributions against the fitted gamma distributions, and statistically with one-sample Kolmogorov–Smirnov (KS) tests (Figure 2d-f). In humans, no individual KS test reached significance (mean KS = 0.030, mean $P = 0.558$ with 0/22 rejected at $\alpha = 0.05$), indicating that the gamma renewal model provided an excellent fit. In macaques, the KS tests were significant for both animals (monkey C: k = 14.7, KS = 0.025, $P < 0.001$; monkey G: k = 10.2, KS = 0.014, $P = 0.014$), and the marmoset likewise showed a significant deviation (k = 18.5, KS = 0.068, $P < 0.001$). To characterize the nature of these deviations, we compared the empirical variance, kurtosis, and quantile ranges of the rescaled ISIs against the predictions of the fitted gamma. In monkey C (k = 14.7), the central body was more concentrated than predicted ((interquartile range, IQR, 9% narrower than the fitted gamma IQR) while the extreme tails were wider (1st–99th percentile range 7% wider), and kurtosis exceeded the fitted gamma prediction (empirical: 1.54, fitted: 0.41), indicating a leptokurtic distribution, that is, a distribution with a sharper peak and heavier tails than the fitted gamma. In monkey G (k = 10.2), the KS rejection was marginal ($P = 0.014$) and quantile ranges matched the fitted gamma closely (all ratios 0.97–1.02). The significant KS test is therefore likely due to slight deviations from a pure gamma distribution combined with high statistical power provided by the large number of ISIs. The marmoset showed a similar pattern as monkey C, but even more pronounced (IQR 19% narrower, 1st–99th range 22% wider, excess kurtosis empirical: 3.49, fitted: 0.33). The leptokurtic pattern in monkey C and the marmoset is

consistent with a highly regular core timing process plus few very long ISIs. Those very long ISIs might be due to the saccade-detection procedure missing a few saccades.

Figure 2c shows the distribution of ISIs after rescaling each interval by its local saccade rate, which removes rate variation across trials and isolates the local timing regularity of the saccadic rhythm.

Having established a substantial rhythmic component in primate saccade timing, we next asked how slow rate fluctuations affect spectral measures that have been used to assess saccade rhythmicity (Amit et al., 2017). We simulated stationary and non-stationary gamma renewal processes that had similar local regularity and similar mean rates yet differed with regard to the presence or absence of rate fluctuations over time (Figure 3).

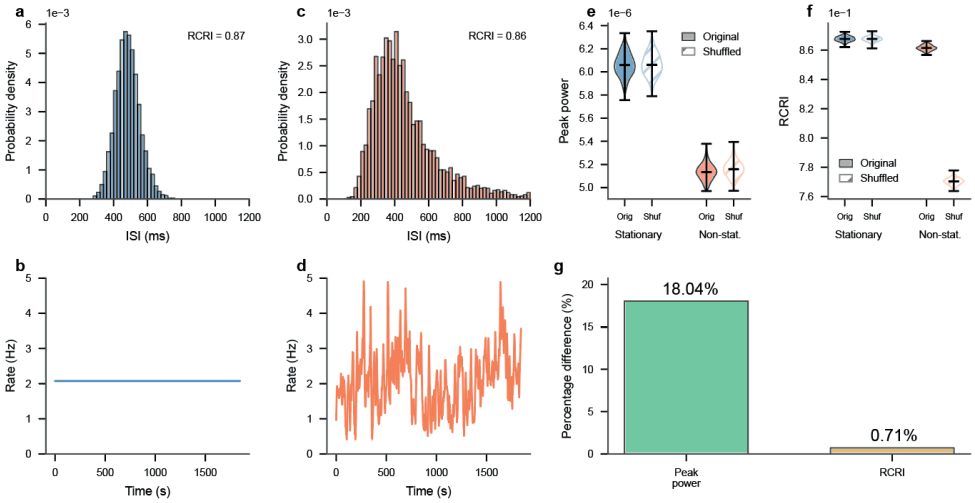

Figure 3. Slow rate fluctuations confound spectral rhythmicity measures.
All panels show the results of simulations of gamma renewal processes with a fixed shape parameter of 44.4. **(a, b)** Example stationary gamma renewal process with fixed rate that is based on the mean of the time series in d. The resulting marginal ISI distribution is narrow and devoid of a heavy tail. Note that the x-axis scale in panel B is in seconds. (**c, d)** Example non-stationary gamma renewal process with an empirical rate time series from a free viewing trial of the MOVIE condition. The resulting marginal ISI distribution is heavy-tailed despite being based on the same shape parameter as panel a and thereby the same local interval regularity. Note that the x-axis scale in panel d is in seconds. **(e)** Peak power (2–6 Hz) across 500 simulations, comparing original and ISI-shuffled spike trains for both stationary and non-stationary gamma renewal processes. **(f)** Rate-corrected rhythmicity index (RCRI) for the same original and ISI-shuffled stationary and non-stationary processes across 500 simulations. Note that the y-axis is highly zoomed in and reveals that the RCRI is highly similar between the stationary and non-stationary process, and that shuffling strongly reduced RCRI as required for shuffling to be a valid control. **(g)** Percent difference between the same stationary and non-stationary processes in peak power (left) and RCRI (right).

Figure 3a,b shows an example stationary gamma renewal process with a shape parameter of 44.4 and a constant rate corresponding to the mean rate of the empirical rate time course shown in Figure 3d. This process produced a narrow ISI distribution, and a rate-corrected rhythmicity index (RCRI, see Methods) of 0.87, reflecting its strong local regularity. Figure 3c,d shows an example non-stationary gamma renewal process with a rate that followed a time course empirically observed during a long free-viewing MOVIE-condition trial. Despite producing a heavy-tailed ISI distribution (Figure 3c), the RCRI was nearly identical to the stationary process (RCRI = 0.86), indicating very similar local regularity once slow rate fluctuations were accounted for.

Repeating this simulation 500 times showed that spectral measures of rhythmicity, such as peak power(Amit et al., 2017), are strongly affected by rate changes (Figure 3e), while the RCRI is hardly affected (Figure 3f, note the highly zoomed-in y-axis). The rate-related difference is 18.04% for peak power, while it is merely 0.71% for the RCRI (Figure 3g).

An ISI-shuffled control within each condition (lighter violins in Figure 3e,f) revealed that peak power was unaffected by shuffling regardless of stationarity, demonstrating that spectral measures alone do not validly capture local temporal regularity in the face of slow rate fluctuations. In contrast, RCRI dropped selectively after shuffling in the non-stationary condition while remaining invariant for the stationary process, confirming its sensitivity to the local serial dependencies imposed by rate modulation.

Together, Figures 2 and 3 show that two properties of the temporal dynamics of saccade generation, namely heavy-tailed ISI distributions and attenuated spectral peaks, can both arise from a locally regular saccade generator operating under a slowly varying rate. These slow rate fluctuations will be characterized in more detail later. In this framework, local rhythmicity of successive intervals is preserved, while slow rate fluctuations shape global interval statistics, such as the marginal ISI distribution, and introduce higher-order temporal structure that degrades standard spectral measures of rhythmicity. Having disentangled these rhythmic and rate-dependent components, we next used the RCRI to isolate the rhythmic component and compare it across tasks and species.

Saccade rate explains across-task differences in rhythmicity

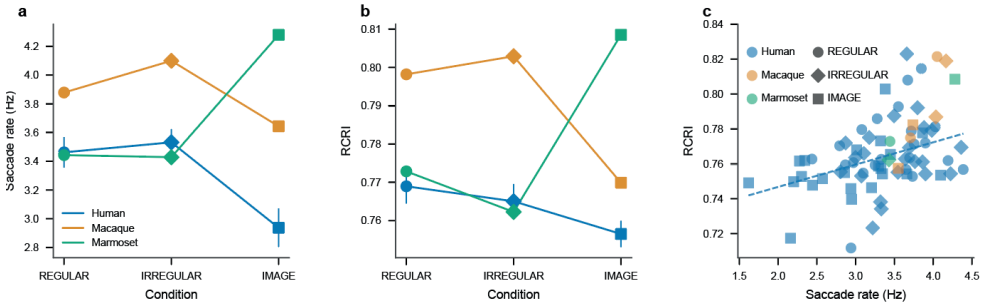

Figure 4. Saccade rates and RCRI across tasks and species.
(a) Mean saccade rate across task conditions (REGULAR, IRREGULAR, IMAGE) for humans, macaques, and marmoset. Task context robustly modulates saccade rate in all species. **(b)** RCRI across the same conditions, showing condition-dependent changes that closely mirror saccade rates. **(c)** Relationship between saccade rate and RCRI for human data (blue symbols). Each symbol corresponds to the combination of a human participant and a task condition. Human data show a significant positive correlation (dashed line, based on human data only). Macaque and marmoset data were not included in the correlation analysis yet fall along the same trend. Importantly, RCRI is invariant to uniform scaling of ISIs (see Methods), meaning the positive correlation between rate and rhythmicity cannot arise from a trivial dependence of the RCRI metric on ISI duration.

So far, we have established that slow rate fluctuations shape ISI distributions and degrade spectral measures of saccade rhythmicity, and that the RCRI provides a robust estimate of local rhythmicity that is nearly immune to slow rate variability. Equipped with this metric, we investigated whether there is a genuine relationship between RCRI and saccade rates, across tasks and species, i.e., a relationship that is not due to a trivial rate confound of the metric but that reflects a true physiological rate dependence. We calculated RCRI for humans, macaques, and marmoset performing a visual foraging task in three conditions: a regular hexagonal grid (REGULAR), a spatially jittered version of the same grid (IRREGULAR), and a natural-image condition (IMAGE). This task design produced various inputs to the visual system by systematically varying spatial periodicity between the REGULAR and IRREGULAR conditions. To further enhance ecological validity, we added the naturalistic IMAGE condition.

Figure 4a shows average saccade rates across conditions for each species. In all species, saccade rates varied substantially between conditions, with the IMAGE condition generally eliciting lower saccade rates in humans and macaques, yet higher rates in the marmoset. Thus, task context robustly modulated saccade rate in all three species.

Figure 4b shows the corresponding RCRI values. Within each species, RCRI varied systematically with condition in a manner closely mirroring changes in saccade rate. Conditions associated

with higher saccade rates also exhibited higher RCRI, whereas conditions with lower saccade rates showed reduced RCRI. This shared modulation suggests a genuine link between saccade rate and local saccade rhythmicity.

To quantify this relationship, we examined the association between saccade rate and RCRI. Figure 4c shows this relationship for human participants, with each dot corresponding to one of the tasks in one participant. This revealed a positive correlation between saccade rate and RCRI ($r = 0.39$, $P = 0.001$). The corresponding data points from the two macaques and the marmoset fall into the human point cloud, descriptively pointing towards a similar relationship across species despite differences in absolute rates.

To disentangle the effects of saccade rate and task condition, we fit a linear model predicting RCRI from average saccade rate and condition. In humans, saccade rate was a significant predictor of RCRI ($\beta = 0.0114$ per Hz, $P = 0.014$), whereas neither the IRREGULAR nor IMAGE condition differed from REGULAR once rate was accounted for. Thus, after controlling for saccade rate, task condition no longer explained additional variance in rhythmicity. Similar patterns were observed in macaques, where rate showed a strong positive effect on RCRI ($\beta = 0.146$ per Hz, $P = 0.026$), while condition effects were not significant. The marmoset data followed the same trend descriptively, although formal inference was not possible due to the limited sample size.

Together, these results demonstrate differences in RCRI across tasks and species. Intriguingly, while the RCRI metric corrects for slow rate fluctuations, the RCRI differences could be accounted for by differences in saccade rate. This suggests a genuine relation between saccade rate and saccade rhythmicity.

Saccade rhythmicity as function of saccade rate

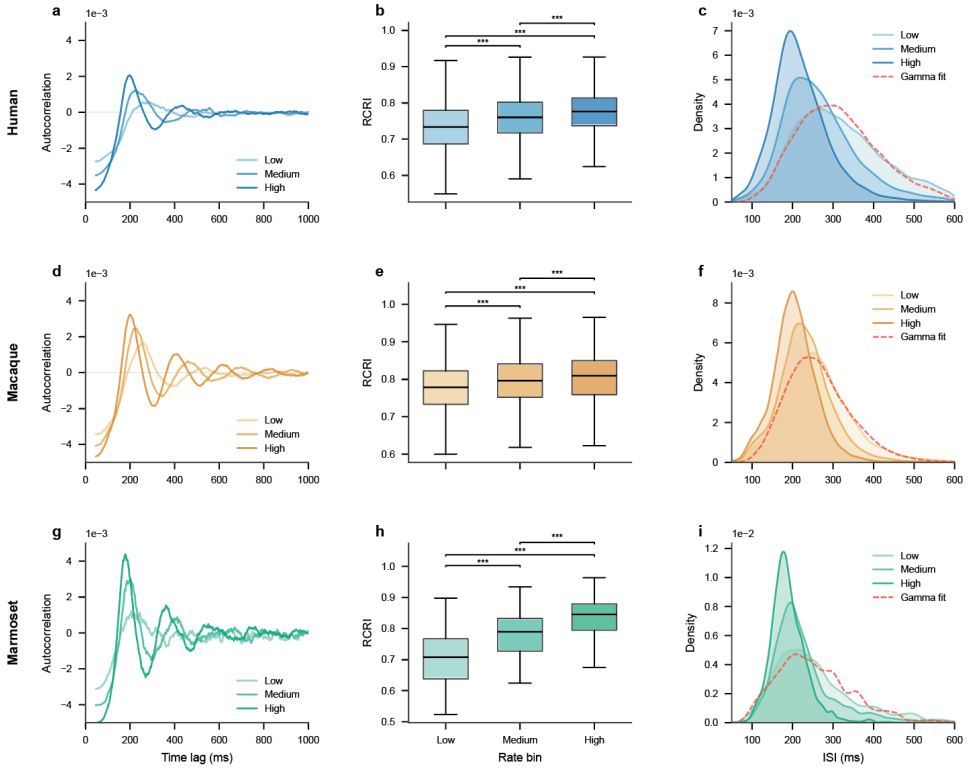

Figure 5. Rhythmic structure strengthens monotonically with saccade rate.

(a, d, g) Autocorrelation functions of saccades for trials binned into low, medium, and high saccade-rate tertiles, shown separately for humans, macaques, and marmoset. Higher-rate trials exhibited more pronounced oscillatory side peaks. **(b, e, h)** RCRI across rate bins for each species. RCRI increases monotonically from low- to high-rate bins. **(c, f, i)** ISI distributions for the same rate bins. Increasing saccade rate is associated with shorter ISIs and narrower ISI distributions. We repeated the analysis of Figure 2 for the lowest saccade rate bin and plotted the ISI distributions approximated via doubly stochastic gamma fit as red dashed lines. Even for the lowest rate bin, all gamma shape parameters are > 8, indicating substantial symmetry.

As an independent and nonparametric validation of the relationship between saccade rate and rhythmicity, we next binned trials by their mean saccade rate and assessed rhythmicity within these rate-defined trial subsets. For each species separately, trials were assigned to tertiles of saccade rate (low, medium, high) using quantile binning, ensuring equal numbers of trials per bin within species and controlling for baseline rate differences across species (see Figure. 1).

Across humans, macaques, and marmoset, higher-rate trials consistently exhibited stronger rhythmic structure in the autocorrelation function (Figure 5a,d,g). In all species, the high-rate bin showed the clearest side peaks at short lags, while rhythmic structure weakened progressively in the medium- and low-rate bins. Consistent with this pattern, RCRI showed robust, monotonic increases with saccade rate (Figure 5b,e,h). In humans, RCRI differed between all pairs of rate bins (high vs. medium: $t = 7.944$, $P = 2.55 \times 10^{-15}$; high vs. low: $t = 18.082$, $P = 1.36 \times 10^{-68}$; medium vs. low: $t = 11.091$, $P = 5.86 \times 10^{-28}$). The same pattern was observed in macaques (high vs. medium: $W = 4.664$, $P = 3.10 \times 10^{-6}$; high vs. low: $W = 11.140$, $P = 8.01 \times 10^{-29}$; medium vs. low: $W = 6.980$, $P = 2.96 \times 10^{-12}$) and in the marmoset (high vs. medium: $W = 6.739$, $P = 1.59 \times 10^{-11}$; high vs. low: $W = 11.171$, $P = 5.67 \times 10^{-29}$; medium vs. low: $W = 6.793$, $P = 1.10 \times 10^{-11}$). Correspondingly, ISI distributions shifted systematically toward shorter ISIs and narrower distributions with increasing saccade rates (Figure 5c,f,i).

This rate-binned analysis provides converging evidence that saccade rhythmicity increases monotonically with saccade rate for each of the investigated primate species. Importantly, this conclusion is supported simultaneously by spectral analysis such as the structured side peaks in the autocorrelation function, by marginal measures such as systematic shifts in ISI distributions, and finally by the RCRI. Thus, the observed relationship between saccade rate and rhythmicity is not an artifact of a specific metric, modelling assumption, or analysis framework, but reflects a robust property of saccade timing across species.

Slow rate fluctuations induce higher-order temporal dependencies

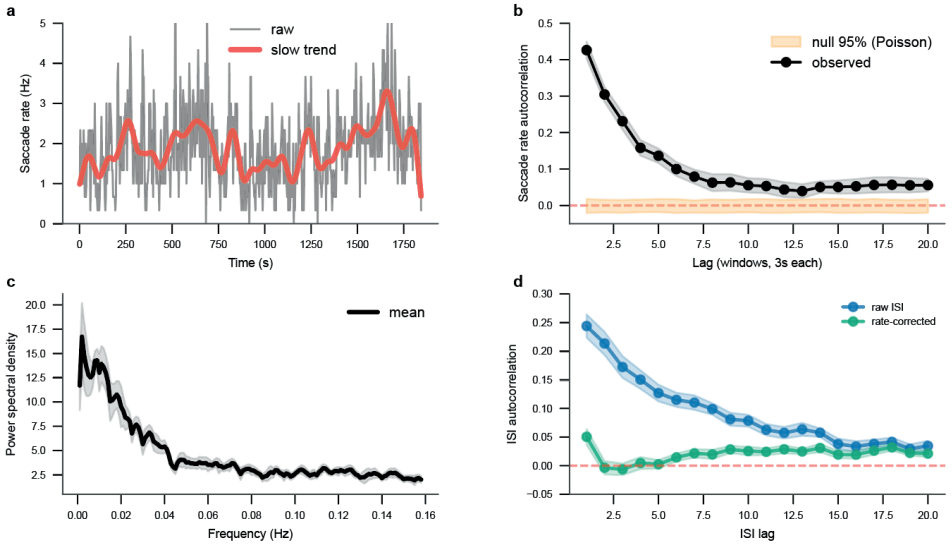

Figure 6. Slow rate fluctuations generate higher-order temporal dependencies.

(a) Example saccade-rate time course from a single free-viewing MOVIE-condition trial, estimated in non-overlapping 3 s windows. Gray line shows raw windowed rates; red line shows a low-frequency trend illustrating slow rate drift over tens to hundreds of seconds. Note that the x-axis scale in panel a is in seconds. **(b)** Serial correlation of windowed saccade rate across trials (black, mean ± SEM across participants), compared to a Monte Carlo null envelope derived from stationary homogeneous Poisson processes (orange, 95% CI). Non-overlapping 3-second windows. **(c)** Power spectral density of windowed saccade rate (mean ± SEM across participants), showing concentration of power at very low frequencies, consistent with slow rate fluctuations. **(d)** Serial correlation of ISIs before (blue, mean ± SEM across participants) and after time-rescaling by the estimated time-varying rate (green, mean ± SEM across participants). Rate correction collapses ISI correlations toward zero across lags.

Having established that saccade rate varies substantially across trials (Figure 1), we next investigated saccade-rate dynamics during free viewing of the MOVIE condition in ≈30 min long uninterrupted trials. Figure 6a shows an example saccade-rate time course from a single free-viewing trial. Although instantaneous rate estimates were noisy, a pronounced component fluctuating on the timescale of tens to hundreds of seconds is evident. We quantified this temporal structure by computing the serial correlation of the windowed rate sequence (Figure 6b, see Methods for the relation between serial correlation and autocorrelation). This rate serial correlation was positive at short lags and decayed gradually with increasing lag, suggesting slow temporal structure. To verify that this effect did not arise trivially from estimating rate in finite windows, we calculated the serial correlation also for a Monte Carlo null model. Surrogate rate time series were generated from stationary homogeneous point processes with no rate drift, conditioned on the observed total number of saccades per trial and analyzed using the

identical windowing and serial correlation procedure. Under this null model, expected rate serial correlation was close to zero at all lags, with a narrow 95% confidence envelope reflecting finite-sample variability (Figure 6b, orange). This null model did not show the empirical serial-correlation decay.

Spectral analysis of the rate time series further supported the presence of slow rate fluctuations. The power spectral density of the windowed rate estimates was dominated by very low frequencies and showed a ~1/f-like decay, with no distinct spectral peak (Figure 6c).

We investigated whether these slow rate fluctuations induced higher-order temporal dependencies in saccade timing. We computed the serial correlation of ISIs excluding intervals overlapping blinks. ISI serial correlation was positive across many lags (Figure 6d, blue), indicating that successive intervals are not independent. These higher-order dependencies extended well beyond adjacent intervals and decayed gradually with lag.

Critically, to test whether these ISI correlations reflect intrinsic structure of the saccade generator or are instead induced by the slow rate variability, we applied a time-rescaling transformation. Each ISI was converted from clock time to a rate-corrected measure by integrating the empirically estimated instantaneous rate over a 3 s analysis-window duration. This yielded the cumulative expected number of events during that interval, accounting for rate fluctuations. If dependencies arose solely from rate drift rather than intrinsic higher-order mechanisms, rescaled intervals should become approximately independent. After transformation, the serial correlation collapsed to values near zero across all lags (Figure 6d, green). This marked reduction demonstrates that the apparent higher-order dependencies in the ISI sequence are largely induced by slow fluctuations in saccade rate, rather than reflecting intrinsic temporal structure of the saccade generator itself. In other words, once rate fluctuations are accounted for, successive saccade intervals behave approximately as a renewal process: each interval depends only on the current rate, not on preceding intervals.

Together, these results show that saccade rates within long trials are non-stationary and exhibit slow, scale-free drifts that are not merely due to noisy ISI variation. This non-stationarity, specifically the slow rate drift, generates temporal structure in the ISI sequence.

Discussion

We explored the temporal dynamics of human, macaque, and marmoset saccades during a variety of different visual stimulation and task conditions. Across the three primate species performing matched free viewing visual foraging tasks, we find that saccade timings can be decomposed into regular rhythmic intervals and non-stationary rate modulation. First, saccade

rates varied widely across trials and drifted slowly within long trials, implying that the empirical heavy-tailed ISI distributions are mixtures over many local regimes rather than signatures of unitary stationary processes. Consistent with this, a gamma-renewal model with a high shape parameter and time-varying rate reproduced the ISI heavy tail, and ISI distributions became more symmetric and narrower after normalizing for slow rate fluctuations. Second, we showed that slow rate fluctuations can strongly distort spectral indices of rhythmicity, attenuating peaks even when local regularity is unchanged and thereby confounding previously used peak-power based quantifications and interpretations of saccade rhythmicity. To address this, we introduced the RCRI that is largely invariant to slow rate drifts and then used it to compare saccade rhythmicity across tasks and species. With this metric, across-condition differences in rhythmicity were exclusively explained by differences in saccade rate. Non-parametric tests revealed a monotonic strengthening of rhythmicity with increasing rate for all three primate species. Finally, slow rate dynamics induced long-range serial correlations of ISIs, but these dependencies collapsed under time-rescaling, indicating that the dominant higher-order structure arises from smooth rate drift rather than intrinsic multi-interval dependence of the local generator. Together, these findings reconcile heavy-tailed ISIs and weak spectral peaks with a robust local rhythmic component whose expression is governed by how the oculomotor system modulates rate across time and context.

Several limitations to our findings should be noted. First, for the macaques and the marmoset, the sample sizes were two and one, respectively, such that the inference was limited to the sample rather than to the population (Fries & Maris, 2022). Second, our conclusions are tied to the specific stimulus set and viewing regimes tested here (static arrays, natural images and relatively slow/smooth movie free viewing). Saccade statistics, including higher-order dependencies, may differ for stimuli and/or tasks with more pronounced temporal structure, or for other behavioral goals. Third, although the deep learning-based saccade detection is designed to be highly sensitive, rhythmicity and serial-correlation estimates can, in principle, be influenced by systematic detection biases, particularly for small-amplitude saccades and during free viewing where other eye-movement components such as smooth pursuits are present. Finally, our modelling framework is intentionally descriptive and based on behavioral data. It constrains the structure of the generative process but does not identify the underlying neuronal mechanism, motivating future work that links rate modulation and local rhythmicity to circuit-level dynamics.

The temporal structure of saccade sequences has a long history of empirical characterization and theoretical debate. Early work established that successive saccades cannot generally be initiated arbitrarily close in time. This post-saccadic refractory period imposes an apparent hard floor on the ISI, on the order of 100–200 ms (Hallett & Lightstone, 1976; Hallett & Adams, 1980) dictating the primary mode of the ISI distribution. Mechanistically, it was suggested by Carpenter and colleagues that this mode reflects the time for a threshold to be reached by a linearly

accumulating decision signal with a stochastically varying rate, producing characteristically right-skewed reaction times (or ISI distributions) as a natural consequence of this variability (Carpenter & Williams, 1995; Noorani & Carpenter, 2016). Sumner and colleagues extended this view within neuronal field models to show that competitive fixation-saccade dynamics further shape the tail (Bompas & Sumner, 2011). Together, these frameworks converge on the view that the heavy-tailed ISI distribution is primarily a product of stochastic accumulation-to-threshold dynamics gated by a post-saccadic recovery period.

The most direct empirical evidence that saccade timing can nonetheless express a genuine local rhythmic component comes from the work of Otero-Millan and colleagues. Examining ISI distributions across saccade amplitude classes during both fixation and free exploration Otero-Millan et al. (2008) showed that the statistics of successive events are consistent with a single common saccadic generator whose inter-event distributions share the same ex-Gaussian form regardless of saccade amplitude, with a mean of the Gaussian component near 150 ms and a heavy tail reflecting occasional prolonged intervals. The case for an intrinsically rhythmic generator becomes compelling in patients with square-wave jerks, spontaneous conjugate saccadic intrusions that interrupt fixation in paired saccades away from and back to a fixation location. Otero-Millan et al. (Otero-Millan et al., 2011; Otero-Millan, Schneider, et al., 2013) documented that the interval between these paired saccades converges tightly on approximately 200 ms, consistent with square waves observed in healthy individuals (Hafed & Clark, 2002; Hafed & Ignashchenkova, 2013). Rather than being merely an effect of pathology, these near-constant inter-saccadic periods could be understood as revealing, in unusually clean form, the local rhythmic kernel that the oculomotor system's intrinsic recovery dynamics impose. This kernel might normally be obscured in healthy free viewing by the superposition of rate variability, higher cognitive influences, or other sources of variability. This directly supports our finding that a robust rhythmic component is present in healthy primate saccade sequences once slow rate fluctuations are accounted for. Even when modelling how saccade generation is interrupted by sensory transients (Buonocore & Hafed, 2023), successive saccade generation was an outcome of repetitive rise-to-threshold waves that could be reset by exogenous stimuli (Hafed & Ignashchenkova, 2013), again utilizing the concept of saccade-system rhythmicity.

Despite this rich tradition of characterizing saccade timing, none of these accounts set out to quantify saccade rhythmicity as such. To our knowledge, the only study that has considered this question directly is Amit et al., (2017), who explicitly asked whether free-viewing saccade sequences are better described by a self-paced generator or by a fixed-frequency oscillator, i.e., a central pattern generator. To answer this question, Amit et al. quantified spectral modulation strength, complemented by point-process analyses of autocorrelograms and hazard functions, and used ISI shuffling to test for dependencies beyond the immediately preceding saccade. In their data, hazard functions were dominated by post-saccadic inhibition without sustained periodic modulation, autocorrelograms showed at most short-lived damped structure, and

shuffling left the oscillatory-modulation index largely unchanged. They concluded that a 3–6 Hz spectral peak is parsimoniously explained by first-order dynamics such as a post-saccadic refractory period, without invoking an autonomous pacemaker.

Our results extend this framework by showing that its core diagnostics can be confounded when saccade rate is non-stationary. Amit et al.'s generative models assume a constant base rate, and the ISI shuffle test is naturally interpreted in this stationary setting. In our data, saccade rate varies strongly across trials and drifts within them, so the empirical ISI distribution becomes a marginal distribution that pools intervals sampled at different prevailing rates. Rate modulation can therefore affect global timing statistics, masking spectral signatures typically attributed to rhythmicity. The same logic applies to the hazard function: locally present periodic modulation can be washed out by the mixture, yielding the flat or monotonically rising profiles that Amit et al. interpreted as evidence against an oscillatory process. Our simulations confirm that shuffling leaves peak power essentially unchanged not only for stationary renewal data but also for non-stationary, rate-modulated data (Figure 3e), so a null shuffle effect does not uniquely support the inference that only first-order dependencies are present. By contrast, the RCRI is shuffle-invariant for stationary renewal data but is disrupted by shuffling for non-stationary sequences (Figure 3f), because shuffling breaks the alignment between successive intervals and the slowly varying rate regime that generated them. These results do not contradict Amit et al.'s conclusion that saccade timing is poorly captured by a rigid, continuously phase-advancing oscillator, but extend their analysis to the regime that dominates natural viewing. When rate variability is accounted for, a local rhythmic component becomes evident, while much of the long-range ISI structure is explained by slow rate modulation.

Our descriptive model decomposes saccade timing into a local gamma-renewal generator and slow, empirically observed non-stationary rate modulation. Together, these two components account for the heavy-tailed marginal ISI distribution and for long-range serial dependencies that collapse under time-rescaling. The natural mechanistic question is which neuronal architecture could implement this decomposition. A parsimonious candidate is a drift-to-threshold mechanism. In this framework, after each post-saccadic reset, a decision variable accumulates toward a fixed threshold. At threshold crossing, a saccade is triggered, followed by a refractory period before the next build-up begins. Each cycle is analogous to one period of an oscillator, but the frequency is not fixed, it varies with the drift rate and the threshold level. This mechanism directly explains our central empirical finding that RCRI increases with mean saccade rate: If the noise of the accumulation process is approximately stable, increasing the mean drift compresses the time-to-threshold distribution and reduces relative dispersion, producing a more regular local sequence for shorter intervals (Figure 7).

This mechanistic explanation has strong neurophysiological grounding. In the frontal eye fields (FEF), movement-related neurons reach a relatively fixed discharge level at saccade onset, with latency variability explained by variability in the growth rate toward threshold (Hanes & Schall, 1996). Analogous threshold-crossing dynamics have been reported in the superior colliculus (Dorris et al., 1997). Downstream, the brainstem burst generator provides an explicit gating and refractory mechanism: omnipause neurons tonically inhibit premotor burst neurons during fixation and pause around saccades, naturally implementing a post-saccadic reset epoch and separating fast, local trigger timing from slower control signals that modulate readiness. Together, these circuits map onto our behavioral decomposition of a locally regular threshold-and-reset generator whose instantaneous rate is set by external drive.

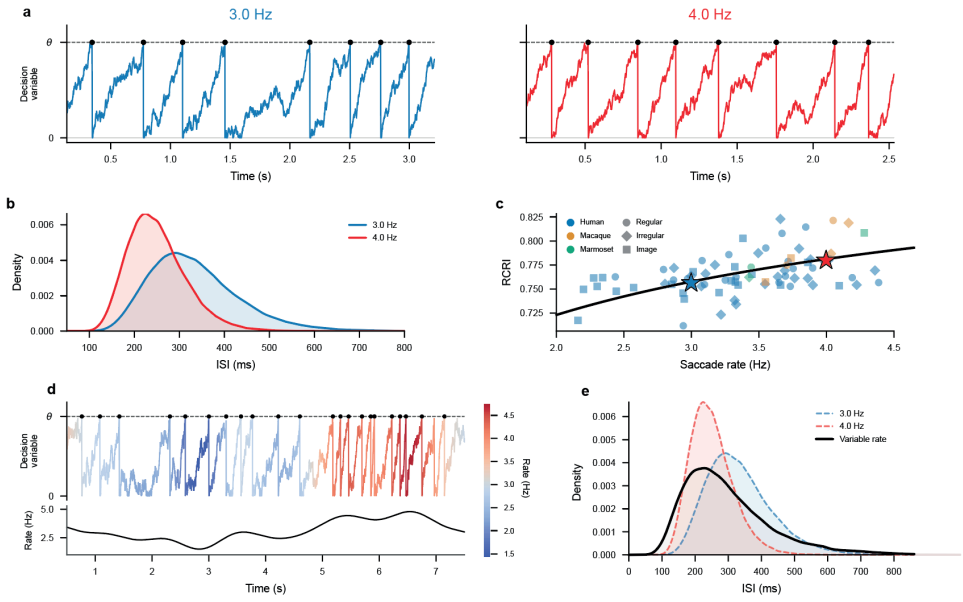

Figure 7. A drift-to-threshold model predicts the observed coupling between saccade rate and rhythmicity.

(a) Example accumulator traces from a drift-to-threshold model at two drift rates (3.0 Hz, blue; 4.0 Hz, red). A decision variable accumulates with constant drift and Gaussian noise toward a fixed threshold (dashed line), triggering a saccade upon crossing. Higher drift produces faster saccade rates. **(b)** Inter-saccade interval (ISI) distributions for the two example rates. Higher drift yields a narrower, more regular ISI distribution, corresponding to higher rhythmicity. **(c)** Rate–RCRI relationship from Figure 4 predicted by the drift-to-threshold model (black line; single free parameter σ fitted to empirical data) overlaid on empirical data from humans (light blue), macaques (orange), and marmoset (green) across conditions. Stars indicate the two example rates from panels a and b. The model captures the monotonic increase of RCRI with saccade rate observed across species and conditions. **(d)** Accumulator trace driven by a slowly fluctuating rate signal, with color indicating instantaneous rate. The rate time course (bottom) varies smoothly between 2.0 and 4.5 Hz, producing epochs of faster and slower saccade generation. **(e)** ISI distribution from the variable-rate simulation (black) compared to the two constant-rate distributions from panel b (dashed). The variable-rate distribution emerges as a mixture of the underlying constant-rate distributions, demonstrating how slow rate fluctuations broaden the overall ISI distribution while preserving local rhythmicity.

The classic Linear Approach to Threshold with Ergodic Rate (LATER) framework assumes that the drift rate is drawn independently on each saccade cycle (or trial, in the classic LATER model), predicting no long-range serial structure in ISIs beyond the refractory period. The slow non-stationarity we observe violates this assumption: the rate does not reset independently after each saccade but instead shows serial correlation structure and therefore drifts continuously, such that successive cycles share a similar local rate context. This can be thought of as the signature of a system operating under slowly drifting control variables such as baseline excitability, urgency, exploration policy, tonic fixation drive, and neuromodulatory state, all of which change on timescales longer than a single ISI. Changes in these control variables might well correspond to different cognitive states, like levels of arousal, reward expectancy and resulting motivation (Bray & Carpenter, 2015; Milstein & Dorris, 2007, 2011; Prsa et al., 2010).

Finally, this framework makes concrete predictions about how rate and local rhythmicity can dissociate. In drift-to-threshold terms, mean rate varies via the mean drift or the threshold level, whereas rhythmicity at matched rate is determined by the variability in drift, reset depth, or refractory gating. A central implication of this is that saccade timing during natural viewing can be viewed as the superposition of at least two timescales: a locally regular generator and slower, non-stationary rate modulation that reshapes global statistics. These timescales might in future studies be investigated during perturbation at various nodes of the saccade control and generation system, thereby identifying the neuronal bases behind the distinct behavioral components. At the neuronal level, saccade generation and attentional control are tightly intertwined. Thereby, a deeper understanding of the neuronal sources of saccade rhythmicity will also provide a link to the rich literature on rhythmicity in attentional and perceptual processes (Bosman et al., 2012; Fries, 2023; Hafed, 2013; Landau et al., 2015).

Acknowledgements

The authors would like to thank Ece Nur Uygur for help with the human data acquisition. This work was supported by DFG (FOR 1847 FR2557/2-1, FR2557/7-1, TH 425/17-1).



5

Chapter 5

FEF Inhibition Increases Saccade Rhythmicity

This chapter is based on the following manuscript:

Näher, T., Breuer, F., Lueckel, M., Bergmann, T. O., & Fries, P. (2026). FEF Inhibition Increases Saccade Rhythmicity .

[Manuscript in preparation]

Abstract

Primates use saccadic eye movements to sample their visual environment. This sampling has temporal structure at two time scales: at a slow time scale of several seconds, saccade rate fluctuates; on top of this, consecutive saccades follow a clear rhythm. Several lines of evidence suggest that saccade timing can be adjusted by top-down cortical control according to momentary perceptual and cognitive demands, which in effect reduces ongoing saccadic sampling rhythmicity. Here, we causally test this hypothesis by applying inhibitory continuous theta-burst stimulation to the right frontal eye field (FEF) in human participants performing free-viewing visual search, with vertex stimulation as control. FEF inhibition indeed increased saccade rhythmicity; saccade rate remained unaffected. Intermittently recorded resting-state EEG, projected to source level, showed that inhibition of the FEF led to a locally increased $1/f^n$ spectral slope, consistent with local cortical inhibition, and reduced beta power. Functional connectivity analysis further revealed attenuated alpha-band connectivity between right FEF and the dorsal attention network. All effects were strongest shortly after FEF inhibition and declined thereafter, matching the expected time course of the intervention. These results demonstrate that the FEF causally shapes the rhythmicity of saccade sequences: when cortical control is reduced, the saccade generator produces more rhythmic output. Under normal conditions, the FEF provides top-down modulation of saccade timing in service of flexible, context-dependent visual exploration, at the cost of rhythmicity.

Introduction

Saccadic eye movements are the primary means by which primates sample their visual environment. Similar active sampling behaviors in rodents are whisking and sniffing. While whisking and sniffing are rhythmic (Kleinfeld et al., 2006), the degree to which primate saccades show comparable rhythmicity has been debated (Amit et al., 2017). Inter-saccadic interval (ISI) distributions are rightward skewed and can be explained by stochastic accumulation-to-threshold dynamics without invoking a rhythmic generator. In a recent study (Näher et al., 2023), we addressed this question using statistical approaches from neuronal spike train analysis and showed that, once fluctuations in saccade rate are accounted for, saccade timing in multiple primate species (humans, macaques, and marmosets) shows strong evidence for inherent rhythmicity. This saccade rhythmicity increases monotonically with saccade rate, a relationship that can be well captured by a drift-to-threshold model, in which a decision variable accumulates toward a fixed firing threshold after each post-saccadic reset.

Here, we investigate whether a key oculomotor area, namely the frontal eye field (FEF), modulates the degree of saccade rhythmicity. In a drift-to-threshold framework, the core rhythmic process maps naturally onto the brainstem burst generator, where so-called omnipause neurons tonically gate premotor burst neurons and contribute to a refractory period after each saccade (Sparks, 2002). These brainstem nuclei and the superior colliculus receive dense projections from the FEF (Segraves & Goldberg, 1987; Stanton et al., 1988). The FEF is well established as a source of top-down signals that shape spatial priority maps and guide saccade target selection (Moore & Armstrong, 2003; Thompson et al., 2005). Given that these projections carry spatial and motor planning signals into the saccade generator, it is conceivable that the FEF activity also affects the rhythmicity of saccade generation. Such an effect on saccade rhythmicity could be a byproduct of variable-latency cognitive processes that modulate the rise to threshold or extend individual fixation epochs. Such a putative link between cortical influence and saccade rhythmicity is supported by two patterns observed in Näher et al., (2023): (1) Across species, saccade rhythmicity was highest in the marmoset, intermediate in macaques, and lowest in humans, consistent with the notion that a larger cortical influence on the saccade generator, present in higher-developed species, reduces saccade rhythmicity. (2) Within each species, rhythmicity was higher for simple dot arrays than for natural images, suggesting that greater demands on top-down visual analysis introduce additional temporal variability into the saccade generator.

Here, we tested the FEF's involvement in saccade timing causally. We applied an offline, inhibitory continuous theta-burst stimulation (cTBS) protocol over the right FEF in human participants performing a visual search task. We targeted the right FEF because, unlike the left FEF, it exerts bilateral control over visuospatial attention. This has been demonstrated both behaviorally (Duecker et al., 2013), via neuroimaging (Ruff et al., 2009), electrophysiologically

(Marshall et al., 2015), and using dual-coil TMS (Silvanto et al., 2006) making right FEF inhibition more likely to affect saccade rhythmicity globally rather than in a single hemifield. cTBS is well established as a tool for producing transient cortical inhibition, and its application to the FEF has been shown to disrupt FEF-dependent functions including top-down attentional control (Duecker et al., 2013; Hung et al., 2011), overt gaze deployment during free viewing (Cazzoli et al., 2015), and the FEF's control over anticipatory alpha oscillations in visual cortex (Capotosto et al., 2009; Marshall et al., 2015).

We leveraged this well-characterized intervention to test a specific prediction: if cortical control from the FEF, e.g. through modulations of individual fixation durations, can override a subcortically generated rhythmic saccade sequence, then cTBS-induced inhibition of the FEF should lead to increased saccade rhythmicity. We find that inhibition of the FEF indeed increased saccade rhythmicity on two independent rate-corrected metrics derived from saccade time series. At the neuronal level, cTBS reduced beta power and steepened the $1/f^n$ spectral slope over the right FEF, consistent with local cortical inhibition, as expected following cTBS. Seed-based orthogonalized envelope correlations further revealed reduced alpha-band connectivity between the right FEF and the dorsal attention network after inhibition, linking diminished top-down coupling to the observed increase in saccade regularity. While the FEF is well established as a key node for spatial attention and saccade target selection, its role in the temporal patterning of saccade sequences has received much less attention. Our results establish that the FEF's influence on saccade generation extends beyond spatial selection to temporal regulation.

Methods

Participants

We recruited eighteen right-handed participants (12 female; aged 22–31, mean = 26.19, SD = 3.58) with normal vision, no history of neurological or psychiatric disorders, and free of medication for the neurostimulation study. Two participants were excluded due to technical problems with eye tracking and EEG recordings, leaving the total number of participants included in the analysis at sixteen. We recruited an additional 10 participants (5 female; aged 21–32) for a preliminary behavioral study to estimate image-level search difficulty for stimulus assignment (see Stimuli). We used these difficulty estimates to calibrate task parameters in the main experiment. All participants provided written informed consent before participation. The study was conducted at two sites (Mainz and Frankfurt am Main), and therefore the study was approved by both, the local ethics committee of the Landesärztekammer Rheinland-Pfalz and the University Hospital Frankfurt am Main. An independent physician screened all participants for MRI and TMS contraindications before inclusion (Rossi et al., 2021). The MRI measurements and the TMS experiments were conducted at the Neurostimulation Lab of the Neuroimaging Center (NIC), Johannes Gutenberg University Medical Center Mainz.

Experimental design

Each participant completed three sessions in a within-subjects design: one anatomical MRI session and two neurostimulation sessions (INHIB-FEF: continuous theta-burst stimulation (cTBS) applied to the right frontal eye field; CTRL: cTBS applied to vertex). The neurostimulation sessions were separated by at least one week. The CTRL session served as a spatial control, providing approximately matched somatosensory and acoustic stimulation without functionally relevant cortical effects. We counterbalanced session order across participants using a fully crossed design, with half of the participants starting with INHIB-FEF and half with CTRL. Each neurostimulation session comprised four blocks: two before (PRE1, PRE2) and two after (POST1, POST2) cTBS delivery. Each block consisted of a 4-min resting-state fixation period, followed by a visual search task (8 images, ~8 min). The total duration of each session was approximately 4 hours including EEG and TMS preparation. The actual data collection and cTBS lasted for approximately one hour.

Stimuli

We derived stimuli from "Where's Wally?" picture books. We divided each of the 34 full-page images into four equal quadrants and excluded the quadrant containing printed instructions, yielding a final pool of 102 images. Because each full image contains only one target character (Wally), the pool comprised 34 target-present and 68 target-absent quadrants. We converted all images to grayscale and contrast-normalized them using histogram matching (Breen & Williams, 1990).

For the 34 target-present images, we aimed at equating search time across experimental blocks (PRE1, PRE2, POST1, POST2). To this end, the preliminary behavioral study participants (n = 10) searched the target-present images, while we recorded the mean and variance of the search time per image. We then removed the two images with the shortest mean search time, to arrive at 32 images, which we divided into 4 images for each of 8 experimental blocks. The images were distributed over the blocks using a simulated annealing algorithm to match mean and variance of search times as precisely as possible. Those 8 experimental blocks corresponded to four blocks (PRE1, PRE2, POST1, POST2) times two sessions (INHIB-FEF, CTRL). The blocks were assigned to participants and session types according to a Latin-square design (Table 1), assuring that each block appeared equally often at each temporal position relative to cTBS delivery (PRE1, PRE2, POST1, POST2) across participants. Per block, we aimed for 8 minutes of search time. To reach this, we added 4 target-absent images per block.

Task

Each trial began with a 2 s central fixation cross, followed by a single image drawn from the current block in randomized order. We instructed participants to search for the target character in the grayscale images and press a button as soon as they found Wally. If participants did not respond within 60 s, the task advanced to the next image automatically.

Each block contained 8 images (4 target-present, 4 target-absent). Importantly, we did not inform participants about the presence of target-absent images. Instead, we told them that every image contained the target, with some being extremely difficult. Target-absent images therefore always ran to the full 60 s time limit. During debriefing, all participants confirmed that they had believed every image contained the target.

Target-present images sometimes were resolved before the 60 s timeout. To nevertheless match block duration across participants and conditions, we implemented the following padding procedure: once a participant had completed all 8 images in a block, we re-presented target-absent images from that block in sequence until total block duration reached 8 minutes. We chose to re-present previously presented target-absent images rather than novel ones, to keep the total amount of novel images constant. We chose to re-present target-absent rather than target-present images to prevent participants from trivially resolving the re-presented stimuli, thereby maintaining engagement and sustained search behavior throughout the full block. This duration-matching procedure was important given the time-sensitive nature of neurostimulation effects. The data acquired during re-presentation of previously presented target-absent images were excluded from analysis.

Table 1. Latin-Square Assignment of Image Blocks to Experimental Conditions and Temporal Positions.

	INHIB-FEF				**CTRL**				
Participant	***PRE1***	***PRE2***	***POST1***	***POST2***	***PRE1***	***PRE2***	***POST1***	***POST2***	***First session***
P01	B1	B2	B3	B4	B5	B6	B7	B8	INHIB-FEF
P02	B2	B3	B4	B5	B6	B7	B8	B1	INHIB-FEF
P03	B3	B4	B5	B6	B7	B8	B1	B2	INHIB-FEF
P04	B4	B5	B6	B7	B8	B1	B2	B3	INHIB-FEF
P05	B5	B6	B7	B8	B1	B2	B3	B4	INHIB-FEF
P06	B6	B7	B8	B1	B2	B3	B4	B5	INHIB-FEF
P07	B7	B8	B1	B2	B3	B4	B5	B6	INHIB-FEF
P08	B8	B1	B2	B3	B4	B5	B6	B7	INHIB-FEF
P09	B1	B2	B3	B4	B5	B6	B7	B8	CTRL
P10	B2	B3	B4	B5	B6	B7	B8	B1	CTRL
P11	B3	B4	B5	B6	B7	B8	B1	B2	CTRL
P12	B4	B5	B6	B7	B8	B1	B2	B3	CTRL
P13	B5	B6	B7	B8	B1	B2	B3	B4	CTRL
P14	B6	B7	B8	B1	B2	B3	B4	B5	CTRL
P15	B7	B8	B1	B2	B3	B4	B5	B6	CTRL
P16	B8	B1	B2	B3	B4	B5	B6	B7	CTRL

Note. Blocks B1–B8 each contained 4 target-present and 4 target-absent images. INHIB-FEF = inhibitory cTBS to right FEF; CTRL = cTBS to vertex control site. PRE1 and PRE2 = blocks before cTBS; POST1 and POST2 = blocks after cTBS.

MRI acquisition

We acquired a T1-weighted anatomical scan for each participant on a 3T Siemens MAGNETOM Prisma MRI scanner, using a 64-channel Siemens head coil (sequence parameters: TR = 2000 ms, TE = 2.99 ms, flip angle = 8°, 0.8 mm isotropic voxels, 256 × 256 × 167 mm FOV).

TMS coil placement and neuronavigation

Two researchers visually identified the right FEF (rFEF) in each participant's T1 image based on the anatomical landmark of the junction between the inferior part of the superior precentral sulcus and the superior frontal sulcus (Paus, 1996; Taylor et al., 2007), placing the target on the anterior bank of the precentral sulcus. In case the location of an identified target differed between the researchers, the anatomy was jointly re-evaluated and a common location agreed upon. For optimal stimulation of the right FEF, we prospectively personalized TMS coil placement (position and orientation) for each participant based on simulations of the TMS-induced electric field (E-field) with SimNIBS (version 4.0.1) (Thielscher et al., 2015). First, biophysically realistic head models were generated by the CHARM pipeline (Puonti et al., 2020). Next, using a validated model of the MagVenture Cool-B65 TMS coil (Drakaki et al., 2022), E-field simulations were run in a search grid of 9 coil positions × 13 orientations (117 placements in total), centered around the scalp point closest to the identified rFEF coordinate. Finally, we selected the placement that maximized mean E-field strength within a 10 mm spherical ROI centered on the individual rFEF coordinate. The optimal placement still allowed two coil orientations, separated by 180 degrees. From those two coil orientations, we selected the one that induced posterior-to-anterior current flow in rFEF to optimally target the individual anterior bank of the precentral sulcus. We guided coil positioning during experiments using MR-based neuronavigation (Localite TMS Navigator v4.0.0, Germany; Polaris Vega ST optical tracker). During neuronavigation, we also digitized the EEG electrode positions in MRI coordinates for each participant and session, which we used for EEG source reconstruction.

TMS setup

We applied TMS with an actively cooled figure-of-eight coil (MagVenture Cool-B65), which was connected to a TMS stimulator (MagVenture MagPro X100 with MagOption) and held in position by a mechanical arm (MagVenture Super Flex Arm) attached to the treatment chair, in which subjects were seated upright. The head position was stabilized using a vacuum pillow to prevent neck strain.

TMS protocol

Following the preparation and registration of electroencephalography (EEG), electromyography (EMG), and electrooculography (EOG), we determined the individual right primary motor cortex (R-M1) hotspot, at which the coil position produced the most consistent and highest-amplitude motor-evoked potentials (MEPs > 50 μV, 15–50 ms post-TMS pulse) in either of the two contralateral hand-EMG channels (see below). To start, we placed the coil tangential to the scalp

on top of the TMS-compatible EEG cap at a 45° angle to the midsagittal plane over the R-M1. We conducted a repetitive grid search of the area with single pulses (biphasic, posterolateral-to-anteromedial induced current direction in the brain tissue for the second half-wave) at increasing stimulation intensities, starting at 50% maximum stimulator output (MSO), while monitoring the EMG for MEPs. Next, we determined the resting motor threshold (rMT) for each participant by applying 40 single TMS pulses at a jittered inter-stimulus interval of 3–4 s to the individual R-M1 hotspot. We used an adaptive staircase procedure implemented in the BEST toolbox (Hassan et al., 2022), starting at an intensity of 5% MSO above the intensity found to elicit MEPs during the preceding hotspot search. We estimated the rMT as the average of the last 10 trials of the staircase ($\overline{rMT}$ = 55.5, SD = 10.5).

For the INHIB-FEF intervention, we delivered cTBS at 80% rMT to the individually navigated rFEF coordinates: 600 biphasic pulses distributed over 200 pulse triplets with an intra-triplet frequency of 50 Hz, and a triplet frequency of 5 Hz, resulting in a total cTBS duration of 40 s. During stimulation and for an additional minute, participants remained seated, fixating a central dot at eye level on the wall before relocating to sit at the desk with the monitor and eye tracker to resume measurements during the POST blocks.

For CTRL, we delivered the same protocol over the vertex, placing the coil on electrode Cz with the coil handle pointing forwards to reverse the induced current direction.

Eye tracking

During task performance, we recorded gaze position from the dominant eye at 1000 Hz using an EyeLink-1000 (SR Research). Participants sat 56 cm from a monitor with a 120 Hz refresh rate (VPixx Technologies) with their head stabilized by a chinrest. We performed eye tracking calibration before each search block (4 times total) and accepted it when mean gaze error was ≤1° of visual angle.

EEG, EMG, and EOG setup

During rest and task periods, 64-channel EEG was recorded using TMS-compatible sintered Ag/AgCl ring electrodes (Multitrodes-TMS, EasyCap), including channels Fp1, Fpz, Fp2, AF7, AF3, AFz, AF4, AF8, F9, F7, F5, F3, F1, Fz, F2, F4, F6, F8, F10, FT9, FT7, FC5, FC3, FC1, FC2, FC4, FC6, FT8, FT10, T7, C5, C3, C1, Cz, C2, C4, C6, T8, TP9, TP7, CP5, CP3, CP1, CPz, CP2, CP4, CP6, TP8, TP10, P7, P5, P3, P1, Pz, P2, P4, P6, P8, PO7, PO3, PO4, PO8, O1, O2, FCz (reference), and POz (ground) (impedances < 5 kΩ). Surface EMG from two hand muscles (abductor pollicis brevis [APB] and first dorsal interosseous [FDI]) was recorded via disposable adhesive hydrogel electrodes (30 × 24 mm) arranged in a tendon-belly montage. Electrooculography (EOG) was recorded via an additional bipolar electrode montage. EEG, EMG, and EOG signals were passed through an anti-aliasing low-pass filter with a low cut at 250 Hz and digitized in DC mode at a 1 kHz sampling

rate by two 7.2 V battery-powered 40-channel 24-bit amplifiers (NeurOne Tesla with Digital-Out Option, Bittium Biosignals, Finland) and acquired using a NeurOne main unit (software v1.5.0).

Gaze and behavioral analyses

Gaze data preprocessing

We recorded gaze position data monocularly at 1000 Hz. We epoched the data from image onset to the participant's response or the 60 s timeout. We identified blinks from the eye tracker's sample status signal. We then extended blink periods by a ±100 ms padding window, set gaze samples within these padded intervals to NaN, and linearly interpolated them. We then low-pass filtered the interpolated gaze time series using a 4th-order Butterworth filter (100 Hz cutoff) applied in both forward and reverse directions to achieve zero-phase distortion.

We detected saccades using a convolutional neural network (Bellet et al., 2019). We trained the network on manually labeled data from a subset of trials from the present study and removed saccades with amplitudes below 15 pixels or durations shorter than 8 ms. We further removed saccades whose peak velocity deviated from the main sequence relationship. Specifically, we fit a power law to the amplitude-velocity relationship across all saccades in log-log space and excluded saccades with residuals above the 99th percentile. We constructed scan paths from the sequence of fixation positions (gaze coordinates at saccade offsets) and defined inter-saccadic intervals (ISIs) as the difference between two subsequent saccade onsets, excluding intervals that overlapped with blink periods. We discarded ISIs shorter than 50 ms or longer than 650 ms. We computed trial-based saccade rate as the total number of detected saccades divided by trial duration.

Saccade Rhythmicity

To capture temporally local rhythmicity while remaining robust to slow fluctuations in overall saccade rate, we computed a rate-corrected rhythmicity index (RCRI) based on the variability between successive ISIs. We previously introduced RCRI in a study quantifying saccade rhythmicity in three primate species (Näher et al., 2023). We assessed local variability using the CV2 statistic:

$$CV_2\,(i) = \frac{2\ |\ ISI_{i+1} - ISI_i\ |}{ISI_{i+1} + ISI_i}$$

We obtained a single trial-level estimate of local variability by averaging $CV_2\,(i)$ across all successive interval pairs, yielding $CV_2 = \langle CV_2(i) \rangle_i$. We then defined the RCRI as:

$$RCRI = \frac{1}{1 + CV_2}$$

This measure ranges from 0 to 1, with values approaching 1 indicating highly regular, rhythmic inter-saccade interval sequences.

Gamma Renewal Model

To further characterize the intrinsic regularity of saccade timing, we fit a gamma renewal model following the approach described in (Näher et al., 2023). This model decomposes ISI variability into two components: (1) a saccade rate r_i that varies slowly with time, and (2) a point process characterized by a gamma distribution with the empirically observed rate parameter r_i and a fitted shape parameter k. The shape parameter k corresponds to the order of a gamma process, i.e. how many exponential stages are summed.

For each ISI, we estimated the local saccade rate r_i in a 3000 ms non-overlapping sliding window centered at the ISI midpoint $m_i = (t_i + t_{i+1})/2$. We truncated the window at trial boundaries and computed the rate as the number of saccade onsets in the window divided by the actual window duration (in seconds).

We rescaled each ISI by its local rate to express it as the expected number of saccades per interval, removing the influence of slow rate fluctuations:

$$\Lambda_i = ISI_i \times r_i \,/\, 1000$$

Under this model, rate-corrected ISIs follow a Gamma(k, k) distribution with unit mean, where k controls the regularity: higher k indicates more regular timing ($k \to \infty$ deterministic; $k = 1$ exponential/Poisson).

We estimated the shape parameter k per participant, condition, and block via maximum likelihood estimation (MLE). We minimized the negative log-likelihood of the Gamma(k, k) distribution over $k \in [0.5, 50]$ using bounded scalar optimization.

Behavioral Metrics

For each trial, we computed: search time (trial duration in ms), number of fixations, saccade rate (saccades/s), mean saccade amplitude (pixels), total scan path length (sum of amplitudes), mean fixation duration, peak saccade velocity, and spatial spread of fixation locations ($\sqrt{(SD_x^2 + SD_y^2)}$). We analyzed only first-presentation trials to avoid repetition effects.

Statistical design

To isolate cTBS-induced inhibition effects from effects of time during the session and unspecific drifts in cognitive states across blocks, we analyzed all behavioral metrics as change-from-baseline scores. For each participant and condition, we defined the baseline as the mean of the

two pre-stimulation blocks PRE1 and PRE2. We computed the percent change score for each post-stimulation block as:

$$\%\Delta_{metric} = \frac{Post_{metric} - Baseline_{metric}}{Baseline_{metric}} \times 100$$

For each metric and post-stimulation block, we compared change scores between INHIB-FEF and CTRL, across subjects, using Wilcoxon signed-rank tests.

Electrophysiological analyses

EEG preprocessing

We performed all electrophysiological analyses in MNE-Python (Gramfort et al., 2013). We recorded EEG continuously from 64 scalp electrodes plus one EOG channel using a BrainVision system. We digitized individual electrode positions using the Localite TMS Navigator v4.0.0 software and stored them as XML marker files for each participant and session.

We preprocessed the data in the following steps. First, we applied zero-phase IIR filters: a 1 Hz high-pass filter followed by a 50 Hz notch filter. These filters were applied to the continuous data, which included enough data before and after relevant task segments (multiple minutes) to avoid filter edge artifacts. We then constructed individualized electrode montages from the digitized positions and epoched the continuous data into the four resting-state segments per session (PRE1, PRE2, POST1, POST2). Data recorded during the active visual search task were not analyzed. We identified bad channels automatically using a Local Outlier Factor algorithm (threshold = 1.5; MNE-Python) and detected muscle artifacts via z-scored high-frequency power (70–140 Hz band, z-threshold = 3.0). After artifact detection, we applied a 100 Hz low-pass filter (zero-phase IIR), resampled the data to 250 Hz, and visually inspected each block to manually annotate any remaining artifact segments.

To ensure a consistent channel set across blocks within a session, we pooled all bad channels identified in any individual block into a single session-level set. We interpolated these channels using spherical spline interpolation and re-referenced the data to the average of all good channels.

As a final cleaning step, we performed independent component analysis (ICA) on the concatenated blocks using extended infomax decomposition. We automatically classified components using ICLabel (Pion-Tonachini et al., 2019) and marked all components not labeled as "brain" for rejection after visual confirmation. We identified additional ocular artifact components by correlating component time courses with the EOG channel. After visually inspecting all component topographies to confirm that only noise components were selected, we applied the ICA solution to each block separately and performed a final visual inspection of the cleaned data to mark any remaining artifacts for exclusion during further analyses.

EEG source reconstruction

We reconstructed individual cortical surfaces from T1-weighted MRI scans using FreeSurfer (version 7.4.1) and defined source spaces on the white matter surface using an octahedral subdivision (approximately 4,098 vertices per hemisphere). For each participant, we constructed a three-layer boundary element model (BEM) with standard conductivities (scalp: 0.3, skull: 0.006, brain: 0.3 S/m) and computed forward solutions individually per session using the electrode positions co-registered to the anatomical MRI.

We segmented the cleaned resting-state data into 5 s fixed-length epochs with an average-reference projection. We computed data covariances from the baseline epochs using a shrinkage estimator (cross-validated sklearn ShrunkCovariance; grid search over 30 log-spaced shrinkage values from 10^{-4} to 1). Because no dedicated noise recording was available, we used an ad-hoc noise covariance (mne.make_ad_hoc_cov), i.e., a diagonal matrix that assumes uncorrelated sensor noise with an identical standard deviation for all EEG channels. For a single sensor type, this matrix only rescales the data uniformly; the beamformer weights are therefore determined by the data covariance and the forward model.. We projected sensor-level EEG into source space using linearly constrained minimum variance (LCMV) beamforming. To ensure that pre-post differences in source-level metrics reflected genuine changes in neural activity rather than differences in the spatial projection, we estimated the LCMV spatial filter once per session from the two baseline blocks and applied this common filter to all four blocks. We configured the beamformer with maximum-power orientation selection and unit-noise-gain weight normalization.

We applied the common LCMV filter to each block's epochs and extracted parcel-level time courses by averaging source activity within each region of the HCP-MMP1 parcellation (Glasser et al., 2016), using the mean-flip method to resolve sign ambiguity.

Source space spectral analysis

After source projection, we computed power spectra for all regions of the dorsal attention network (DAN). The DAN was defined following the Yeo 7-network parcellation (Yeo et al., 2011) mapped onto the HCP-MMP1 atlas (Glasser et al., 2016). This DAN definition comprises a frontal node (FEF, area 6a), the intraparietal sulcus (7Am, 7PC, AIP, LIPd, LIPv, VIP), and visual motion areas (MST, MT, PH, V4t), always bilaterally. For each participant and session, we constructed spatial filters from the two baseline resting-state blocks (PRE1 and PRE2) using a shrinkage-regularized data covariance matrix and an ad-hoc noise covariance. The LCMV filters were then applied to all four blocks (two baseline PRE1 & PRE2, two post-stimulation POST1 & POST2), providing time courses for each vertex, which were then combined over the vertices in the selected HCP-MMP1 parcels and finally over the parcels in the selected DAN nodes. We computed power spectral densities using Welch's method with Hann-tapered 0.5 s segments with 50% overlap and Zero Padding to 2 s for an interpolated 0.5 Hz frequency resolution.

To isolate spectral peaks from the $1/f^n$ background, we estimated the $1/f^n$ background spectral slope using robust regression in log-log space (Manning et al., 2009; Spyropoulos et al., 2018), fitting only frequencies above 5 Hz to avoid distortion from the 1 Hz high-pass filter. We excluded the 47-53 Hz range since we applied a notch filter to remove line noise during preprocessing. The $1/f^n$-background fit was extrapolated across all frequencies, and this extrapolated fit was used to normalize the full power spectrum and thereby obtain the flattened spectrum. Next, we defined individual alpha and beta frequency bands for each participant. To obtain a robust estimate per participant, we combined the data from both the INHIB-FEF and CTRL sessions before defining frequency bands. For this, we scaled each participant's INHIB-FEF and CTRL spectra separately, such that their respective total power summed to one, averaged them over the two sessions, and identified the most prominent spectral peak near 10 Hz (alpha, search range 7–13 Hz) and 20 Hz (beta, search range 15–25 Hz). One participant (P12) was excluded from alpha-band analyses due to the absence of a detectable alpha peak. Individual alpha and beta bands were defined as the peak frequency ± 2 Hz and ± 4 Hz, respectively.

For the right FEF region (FEF + area 6a parcels, averaged), we extracted $1/f^n$-corrected alpha and beta power, and the $1/f^n$ spectral slope from each block. We computed the change-from-baseline scores as percent change, the same way as for the behavioral metrics (see above).

Connectivity analysis

Functional connectivity within the DAN was assessed using symmetrically orthogonalized amplitude envelope correlations (Colclough et al., 2015; Hipp et al., 2012). We extracted area time courses from 5 s epochs for each block, using the same LCMV spatial filters as for the spectral analysis. These time courses were then bandpass-filtered at each participant's individual alpha or beta band. We minimized the influence of spatial leakage on the estimation of the amplitude envelope correlation by using symmetric orthogonalization (Colclough et al., 2015), which applies the inverse square root of the covariance matrix to decorrelate time courses while preserving their relative amplitudes. Finally, we computed pairwise envelope correlations across all HCP-MMP1 parcels.

From the resulting correlation matrices, we assessed connectivity at two levels. First, at the level of the entire DAN, we computed average connectivity change from baseline in percent change and compared INHIB-FEF and CTRL sessions using paired *t*-tests. These results are presented in Figure 3 a-b. Second, at the level of node pairs, we computed connectivity for each pair of the six DAN nodes by averaging pairwise envelope correlations across all parcel pairs corresponding to each node pair (15 unique node pairs). We tested condition differences (INHIB-FEF minus CTRL) using paired *t*-tests and visualized significant edges in Figure 3c.

Results

We tested 16 human participants in a within-subjects design comprising two sessions, during which a TMS-based cortical inhibition protocol was applied either to the right FEF (“INHIB-FEF”) or to the Vertex as a control condition (“CTRL”). Each session comprised two pre- (“PRE”) and two post-TMS (“POST”) blocks (Figure 1a). Each block consisted of a 4-minute resting-state EEG recording and an 8-minute visual search task, during which participants searched for a hidden target in complex scenes ("Where's Wally"). An example scan path is shown in Figure 1b, and corresponding eye position traces in Figure 1c. Search performance did not differ between CTRL and INHIB-FEF ($P > 0.05$). Analyses of saccade rhythmicity are presented in Figure 1d-k and the corresponding electrophysiology results from resting-state EEG in Figures 2 and 3.

rFEF inhibition increases saccade rhythmicity

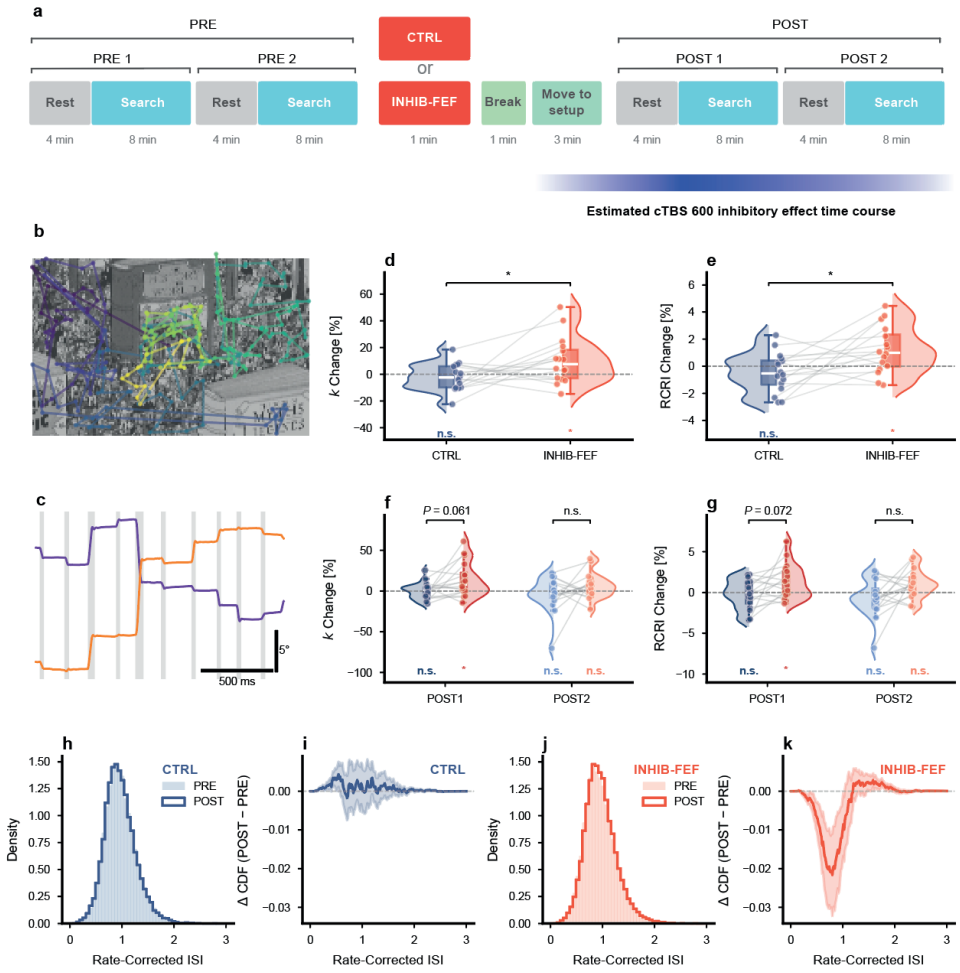

Figure 1. FEF inhibition increases saccade rhythmicity.

(a) Experimental design. Participants completed two sessions in a within-subject crossover design. Each session comprised two pre-stimulation blocks (PRE1 and PRE2, referred to as PRE if pooled) and two post-stimulation blocks (POST1 and POST2, referred to as POST if pooled), each consisting of a 4-minute resting-state EEG recording and an 8-minute visual search task. Continuous inhibitory theta-burst stimulation (cTBS, 600 pulses) was applied between PRE and POST. **(b)** Example scan path from one trial of the visual search task overlaid on the scene. The color of the scan path illustrates the time during the trial. **(c)** Example horizontal (orange) and vertical (purple) eye position traces from one trial. Gray shaded regions indicate detected saccades. **(d)** PRE-to-POST change in gamma shape parameter k for CTRL (blue) and INHIB-FEF (red) sessions. FEF inhibition increased k relative to Vertex (paired t-test between INHIB-FEF and CTRL across participants, $P = 0.038$). Asterisks below each condition indicate within-session change from baseline (paired t-test between POST and PRE across participants); n.s., not significant. **(e)** Same as (d) for the rate-corrected

rhythmicity index (RCRI). RCRI was increased by INHIB-FEF compared to CTRL (P = 0.015). **(f)** Same as d, but separately for POST1 and POST2 (while still pooling PRE1 and PRE2). After multiple comparison correction over POST1 and POST2, k did not change between INHIB-FEF and CTRL. For INHIB-FEF, POST1 compared to PRE showed increases for k. Further POST-PRE comparisons were not significant. **(g)** Same as (f) for RCRI. After multiple comparison correction over POST1 and POST2, RCRI did not change between INHIB-FEF and CTRL. For INHIB-FEF, POST1 compared to PRE showed increases for RCRI. Further POST-PRE comparisons were not significant. **(h)** Rate-corrected ISI distributions for PRE (filled) and POST (outline) during CTRL sessions. **(i)** Difference in cumulative density functions (Δ CDF) (POST − PRE) for CTRL sessions. Shading indicates ± SEM across participants. **(j)** Same as (h) for INHIB-FEF sessions. Note the narrowing of the distribution after cTBS. **(k)** Same as (i) for INHIB-FEF sessions. The negative deflection indicates that fewer saccades fell into the short-interval range after INHIB-FEF, consistent with increased saccade rhythmicity. In all panels, individual participants are shown as dots connected by lines. Box plots show the median and interquartile range.
*$P < 0.05$.

Based on previous work (Näher et al., 2023), we reasoned that the FEF might top-down introduce temporal variability into the saccade generation process, such that inhibiting the FEF with cTBS should increase saccade rhythmicity. To test this, we quantified saccade rhythmicity using two metrics that capture the regularity of saccade timing. The first metric is the shape parameter *k* that we derive from a rate-corrected Inter-Saccadic Interval (ISI) distribution (Näher et al., 2023; see also Methods). This parameter measures the regularity of the ISI distribution: $k = 1$ corresponds to a memoryless (Poisson) process; $k > 1$ indicates increasingly regular timing, with k approaching infinity for perfect, clock-like rhythmicity. The second metric is the rate-corrected rhythmicity index (RCRI) that quantifies the local regularity of consecutive saccade intervals (see Methods). Critically, both metrics are corrected for overall saccade rate. Without this correction, natural fluctuations in saccade rate, potentially arising from changes in arousal or task requirements, inflate the saccade timing variability, distorting estimates of true saccade rhythmicity (Näher et al., 2023). Rate correction ensures that any observed effects of INHIB-FEF reflect genuine changes in saccade rhythmicity. We assessed these effects as change from baseline (POST minus PRE) and compared INHIB-FEF and CTRL sessions using paired *t*-tests across participants.

INHIB-FEF increased saccade rhythmicity relative to CTRL on both measures. The gamma shape *k* increased after INHIB-FEF compared to CTRL ($t(15) = 2.27$, $P = 0.038$, d = 0.52; Figure 1d). Comparisons between PRE and POST showed that rhythmicity increased following INHIB-FEF ($t(15) = 2.37$, P = 0.032; Figure 1d right), whereas it remained unchanged following CTRL ($t(15) = -0.76$, $P = 0.462$; Figure 1d left). Similarly, the RCRI increased after INHIB-FEF compared to CTRL ($t(15) = 2.76$, $P = 0.015$, d = 0.58; Figure 1e), and comparisons between PRE and POST indicated that RCRI increased following INHIB-FEF ($t(15) = 2.76$, $P = 0.015$; Figure 1e right) but not following CTRL ($t(15) = -1.27$, $P = 0.225$; Figure 1e left). Saccade rate itself did not change from PRE to POST in either INHIB-FEF ($P > 0.05$) or CTRL ($P > 0.05$), ruling out rate as a confound. Importantly, vertex stimulation approximately matches the auditory and somatosensory co-

stimulation of active cTBS, controlling for nonspecific sensory and placebo effects. As an active spatial control, it additionally demonstrates that the effect on saccade rhythmicity is specific to FEF inhibition rather than a general consequence of cortical cTBS.

When we compared the PRE-to-POST changes between INHIB-FEF and CTRL, separately for PRE-to-POST1 and PRE-to-POST2, effects did not reach significance after multiple comparison correction over the two tests (Figure 1f, g). However, both k and RCRI showed trends for both POST1 and POST2 consistent with increases after INHIB-FEF. For POST1, the corresponding *P*-values were close to significance ($P = 0.061$ for k, and $P = 0.072$ for RCRI), and the PRE-to-POST comparisons after INHIB-FEF reached significance. This pattern is consistent with the expected transient inhibition of the FEF by cTBS during POST1 and a return to baseline during POST2.

To visualize how INHIB-FEF reshaped saccade timing, we inspected rate-corrected ISI (RC-ISI) distributions, comparing PRE versus POST, separately for INHIB-FEF and CTRL (Figure 1h,j). Because differences in the density plots are difficult to discern by eye, we computed the difference of the cumulative density functions between PRE and POST, which is more sensitive to subtle distributional shifts (Figure 1i,k). The CTRL CDF difference fluctuated around zero without systematic structure (Figure 1i). In contrast, INHIB-FEF showed a pronounced negative deflection just below the mean RC-ISI and a smaller positive deflection above the mean RC-ISI, consistent with a sharpening of the gamma distribution (Figure 1k).

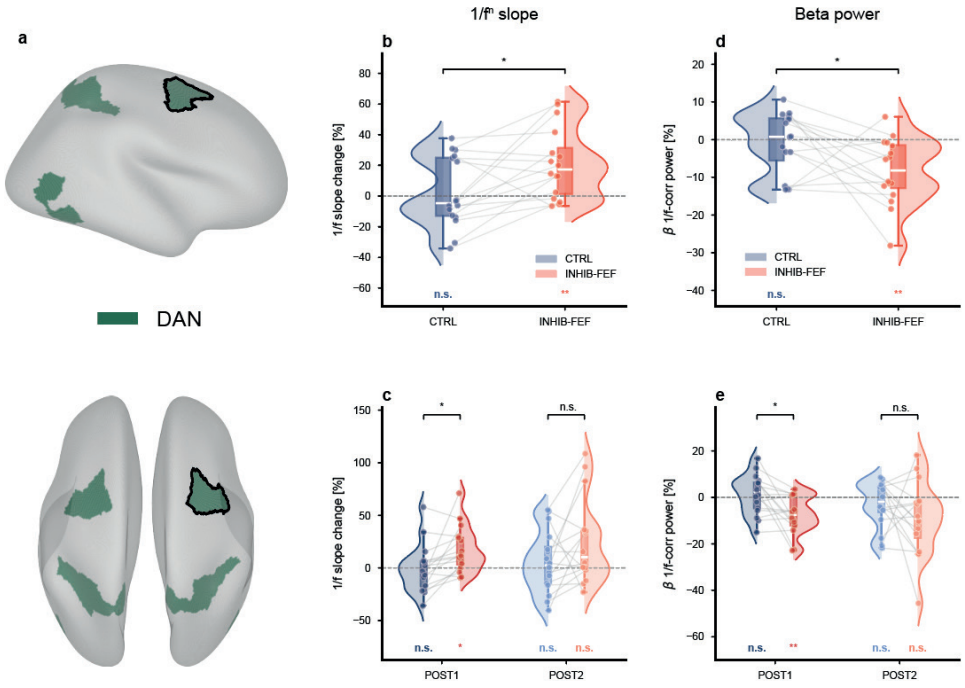

Figure 2. FEF inhibition alters local beta power and 1/fn slope.

(a) Dorsal attention network (DAN, green) regions highlighted on the inflated fsaverage surface. Lateral (top) and dorsal (bottom) views. The right FEF region (FEF + area 6a, HCP-MMP1 atlas) is outlined in black. This region was used for the power analyses in (b-e), and as seed for the connectivity analyses in (f-h). **(b)** PRE-to-POST change for the 1/fn spectral slope (aperiodic component) for CTRL (blue) and INHIB-FEF (red) sessions. FEF inhibition produced a significant 1/fn exponent increase compared to CTRL (P = 0.011). **(c)** Same as (b), but separately for POST1 and POST2. The slope change was significant at POST1 (P = 0.017) but not POST2 (P = 0.179), after correction for multiple comparisons. POST1 compared to PRE showed a significant slope increase specifically for INHIB-FEF, but not for CTRL. **(d)** PRE-to-POST change for 1/fn-corrected beta power at individual peak frequencies for CTRL (blue) and INHIB-FEF (red) sessions. FEF inhibition significantly reduced beta power compared to CTRL (P = 0.033). **(e)** Same as (d), but separately for POST1 and POST2. Beta-power reduction was significant at POST1 (P = 0.016) but not POST2 (P = 0.489), after correction for multiple comparisons. POST1 compared to PRE showed a significant beta-power decrease specifically for INHIB-FEF, but not for CTRL.

*P < 0.05, **P < 0.01, n.s.: not significant.

Having established that inhibiting the FEF increases saccade rhythmicity, we next investigated whether cTBS also altered neuronal activity in the FEF and functional connectivity in the associated dorsal attention network (DAN). We obtained source-level signal time courses from the right FEF during the resting-state blocks, using individual anatomical MRIs, EEG electrode co-registration, and source projection through linearly constrained minimum variance spatial filtering (LCMV, see Methods). From the source-level signal time courses, we computed source spectra and estimated the $1/f^n$ (aperiodic) component using robust regression (see Methods). INHIB-FEF produced an increase in the $1/f^n$ spectral exponent *n* from PRE to POST compared to CTRL ($t(15) = 2.90$, $P = 0.011$; Figure 2b), corresponding to a flattening of the power spectrum. Shifts in $1/f^n$ activity have been linked to changes in cortical excitation–inhibition balance and a flatter spectrum following inhibitory stimulation is consistent with a reduction in net cortical excitability in the FEF (Huang et al., 2005) . When we compared the PRE-to-POST changes between INHIB-FEF and CTRL separately for POST1 and POST2, the slope change was significant at POST1 ($P = 0.017$, corrected for multiple comparisons) but not POST2 ($P = 0.179$; Figure 2c). For POST1, the PRE-to-POST comparison after INHIB-FEF reached significance ($P = 0.006$), whereas it did not for CTRL.

Subsequently, we removed the estimated $1/f^n$ components and used the residual spectrum to identify individual alpha and beta bands as the peaks closest to 10 and 20 Hz, respectively. We detected alpha peaks for 15/16 participants and beta peaks for all 16/16 participants. We then extracted power at each participant's individual peak frequencies with a band-width of ±2 Hz for alpha, and ±4 Hz for beta.

INHIB-FEF reduced beta power relative to CTRL ($t(15) = -2.34$, $P = 0.033$; Figure 2d). When we compared the PRE-to-POST changes between INHIB-FEF and CTRL separately for POST1 and POST2, the beta power reduction was significant at POST1 ($P = 0.016$) but not POST2 ($P = 0.489$; Figure 2e). For POST1, the PRE-to-POST comparison after INHIB-FEF reached significance ($P = 0.002$), whereas it did not for CTRL. We observed no effects for alpha power (all $P > 0.05$).

The observation that both the $1/f^n$ component and beta power showed significant effects at POST1 but not anymore at POST2 mirrors the corresponding observations for saccade rhythmicity. This temporal profile with an early effect and a return to baseline is consistent with the expected time course of cTBS effects in general (Wischnewski & Schutter, 2015) and for the FEF in particular (Marshall et al., 2015).

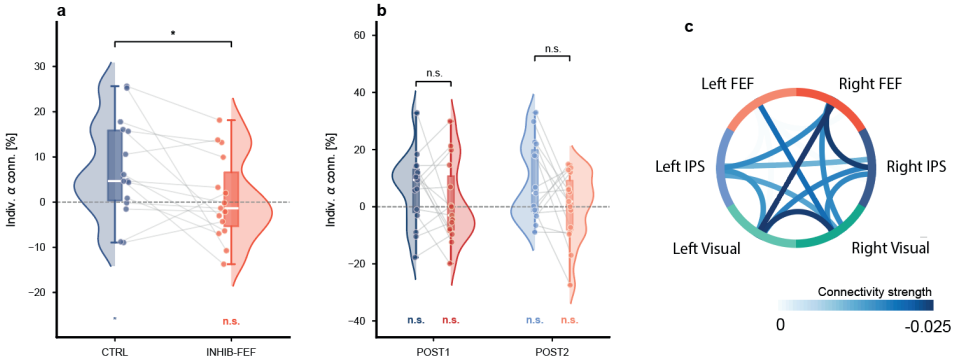

Figure 3. FEF inhibition alters dorsal attention network alpha connectivity.

(a) PRE-to-POST change in DAN alpha envelope connectivity for CTRL (blue) and INHIB-FEF (red). FEF inhibition reduced alpha connectivity between the FEF and the other DAN areas (P = 0.039, averaged over all right FEF connections). **(b)** Same as (a). but separately for POST1 and POST2. Neither comparison survived multiple-comparison correction (POST1: P= 0.538; POST2: P = 0.163). In all panels, individual participants are shown as dots connected by lines. Box plots show the median and interquartile ranges. **(c)** Circle plot showing DAN region-to-region alpha-band orthogonalized envelope connectivity changes (INHIB-FEF minus CTRL). Only significant edges are shown (paired t-test, uncorrected P < 0.05; 9 of 15 pairs). Blue edges indicate reduced connectivity following FEF inhibition. There were no edges with significantly increased connectivity.

*P < 0.05, n.s.: not significant.

Finally, we estimated source-level signal time courses for the other areas in the DAN and investigated the effects of FEF inhibition on symmetric orthogonalized amplitude envelope correlations (Colclough et al., 2015) within the DAN (Figure 2a) at each participant's individual alpha (±2 Hz) and beta (±4 Hz) frequencies. Comparing PRE versus POST, FEF-seeded DAN alpha connectivity was greater following CTRL than INHIB-FEF (t(14) = 2.27, P = 0.039; Figure 3a), indicating that FEF inhibition reduced alpha-band coordination across the attention network. Comparisons between PRE and POST revealed that alpha connectivity increased following CTRL (P = 0.018), whereas it remained unchanged following INHIB-FEF (P = 0.754). When examined separately, neither POST1 (P = 0.538) nor POST2 (P = 0.163) reached significance after correction for multiple comparisons (Figure 3b). We observed no significant effects for beta connectivity (all P > 0.05).

To further characterize the spatial pattern of alpha connectivity changes, we tested all pairwise connections within the DAN (6 nodes, 15 pairs). Nine of 15 pairs showed significantly reduced connectivity following INHIB-FEF compared to CTRL (uncorrected P < 0.05; Figure 3c), with the strongest effects involving connections between the right FEF and bilateral IPS and visual regions.

Discussion

In a previous free-viewing study (Näher et al., 2023) we found that saccade rhythmicity decreases from marmosets to macaques and further to humans, suggesting a rhythmicity-reducing role of cognitive eye-movement control structures like the FEF. The same conclusion was supported by the finding that dot patterns are explored with higher saccade rhythmicity than naturalistic images, whose exploration benefits from variable fixation durations. Based on these observations, we predicted that inhibiting the FEF would increase saccade rhythmicity, reasoning that cortical input to the saccade generator circuits introduces temporal variability into saccade timing.

To investigate whether the FEF modulates the degree of rhythmicity in saccadic eye movements, we applied inhibitory cTBS over the right FEF in 16 human participants performing a free-viewing visual search task. As hypothesized, FEF inhibition increased saccade rhythmicity on two independent rate-corrected metrics, the gamma shape parameter *k* and the rate-corrected rhythmicity index RCRI. This increase occurred without a change in saccade rate. These behavioral effects were accompanied by converging electrophysiological evidence of reduced FEF excitability. Using EEG source reconstruction, we found that cTBS over the right FEF steepened the local $1/f^n$ spectral slope, consistent with reduced cortical excitability, and that it reduced beta power. Seed-based orthogonalized envelope correlations revealed that alpha-band connectivity between the right FEF and the wider DAN was reduced after FEF inhibition relative to control. All behavioral and electrophysiological effects were most pronounced at the first post-intervention block shortly after TMS delivery and attenuated in the second post-intervention block, consistent with the expected time course of cTBS-induced inhibition (Wischnewski & Schutter, 2015). Our results demonstrate that the FEF causally modulates the degree of saccade rhythmicity: when cortical input from the FEF is reduced, the saccade generator produces more rhythmic output. These results establish the FEF as a source of temporal control over saccade generation, complementing its well-characterized role in spatial attention and saccade target selection.

Several limitations of our study should be noted. First, our sample comprised 16 participants. An extension would have, due to the Latin-square counterbalancing design, required integer multiples of this number, i.e., at least a doubling. A larger sample size would have benefitted our cTBS study, given known substantial inter-individual variability and a sizeable non-responder rate (Corp et al., 2020; Hamada et al., 2013). The limited sample size most likely reduced our sensitivity, such that the observed effects should be considered a lower bound. Second, we targeted only the right FEF, chosen for its bilateral influence on visuospatial attention (Duecker et al., 2013; Marshall et al., 2015; Ruff et al., 2009; Silvanto et al., 2006). Whether the left FEF inhibition would produce comparable effects remains open, given evidence for hemispheric asymmetries in FEF function (Duecker & Sack, 2015; Gallotto et al., 2025). Finally, our findings

are tied to a single task regime: free-viewing visual search through complex natural scenes. Saccade rhythmicity may respond differently to FEF perturbation under conditions with stronger top-down demands, such as anti-saccade tasks.

TMS has been widely used to probe the causal role of the FEF in oculomotor control, using both online protocols that disrupt processing on individual trials and offline protocols that reduce cortical excitability across extended periods. These two approaches provide complementary perspectives on how the FEF can shape saccade dynamics. Online paradigms offer temporal precision and reveal the precise time when the FEF contributes to individual saccade programming in each trial. Offline paradigms modulate FEF excitability for sustained periods of several minutes without the need for additional TMS pulses during that period. This enabled us to investigate statistical properties of saccade sequences, without intervening TMS pulses that might have imposed or interfered with internally generated temporal structure.

Online TMS paradigms have explored the effects of disrupting the FEF around the time of saccade programming. When single pulse TMS is delivered over the right FEF during the saccade preparation epoch, saccade latencies increase (Nagel et al., 2008; Thickbroom et al., 1996), and double-pulse TMS over the right FEF impairs target discrimination during visual search within a narrow window of 40 to 80 ms after stimulus onset (O'Shea et al., 2004). The temporal specificity of these effects is consistent with rise-to-threshold dynamics in the FEF and superior colliculus that drive brainstem saccade generators (Dorris et al., 1997; Hanes & Schall, 1996).

Offline inhibitory protocols shift the focus to the sustained modulation of FEF excitability and its effect on various saccade metrics, and have so far yielded notably mixed results. Cameron et al. (2015) found that cTBS over the right FEF reduced anti-saccade amplitudes and increased corrective saccades, but left reaction times and error rates largely intact. Gurel et al. (2018) applied cTBS over the left and right FEF and found no significant effects on either voluntary or reflexive saccade latencies. Gallotto et al. (2025) similarly reported that cTBS over the left or right FEF failed to impair performance on a cued attention task, attributing this null effect to compensatory recovery within the attention network. This preservation of summary metrics such as mean latency, amplitude, and accuracy has been interpreted as evidence for distributed compensation across the frontoparietal oculomotor network, including DLPFC and posterior parietal cortex (Cameron et al., 2020). Importantly, the measures examined in all these studies describe properties of individual saccades. None of them capture the statistical regularity within free-viewing saccade sequences. Note that also in the present study, cTBS of the FEF did not affect search performance.

We suggest that tonic FEF inhibition might have particularly pronounced effects on the statistical properties of saccade sequences: if cTBS reduced FEF activity and thereby cortical input to brainstem saccade generators, the perturbation might primarily alter the variability structure

that governs the regularity of successive intervals. The drift-to-threshold model of Näher et al., (2023) predicts that mean saccade rate and saccade rhythmicity can dissociate: mean saccade rate is mainly determined by the mean drift rate, whereas rhythmicity at a given rate is mainly determined by the variability in drift. If the FEF is a source of drift variability, as the online TMS evidence implies, then reducing FEF excitability should attenuate this variability and increase local rhythmicity without necessarily changing mean rate. This is what we observed in the present study: saccade rate was unchanged after cTBS, while both the gamma shape parameter *k* and the RCRI increased, and the rate-corrected ISI distribution narrowed (Figure 1j,k).

Neurophysiological recordings have linked rhythmic cortical activity to rhythmic attentional sampling. Simultaneous recordings from the FEF and LIP revealed that attention alternates at theta frequency between two states with distinct oscillatory signatures: FEF-dominated beta activity accompanies a sampling state associated with suppressed attentional shifts, whereas LIP-dominated alpha activity accompanies a shifting state associated with attenuated visual processing (Fiebelkorn et al., 2018). Consistent with this, TMS-based perturbations of the human FEF revealed a theta-rhythmic contribution to visual search performance (Dugué et al., 2019). These findings establish that FEF activity is temporally structured in ways that match behavioral sampling rhythms. However, they do not establish where the underlying rhythmicity is generated. The generator might reside in the cortical regions in which it is observed, or it might lie elsewhere, for example in brainstem oculomotor circuits whose intrinsic dynamics impose temporal regularity on saccade release. In the latter case, the rhythmic activity observed in cortex would be entrained by a subcortically generated rhythm rather than being its source.

If the FEF were generating or enhancing saccade rhythmicity, then FEF inhibition should lead to a rhythmicity decrease. Instead, we observe that FEF inhibition leads to a rhythmicity increase. This is consistent with a model in which the underlying rhythmicity is generated outside the FEF. The brainstem saccade generator is a natural candidate, given its role in gating saccade release. FEF's role, on this account, is not to produce or amplify saccade rhythmicity, but to mediate its expression in behavior by flexibly adjusting the durations of individual fixations and the timing of the next saccade. Thereby, the FEF can modulate the rhythm of attentional sampling and thus bias attention to a given location, according to sensory demands and top-down information. Such fixation-by-fixation adjustments introduce variability into the interval structure of the saccade sequence and thereby mask the rhythmicity of the underlying generator. Reducing FEF input attenuates this source of variability and unmasks the rhythmicity that the generator otherwise provides.

The framework developed across (Näher et al., 2023) and the present study generates testable predictions. The drift-to-threshold model specifies that mean rate and rate-matched rhythmicity depend on distinct parameters, so perturbing different nodes should yield distinct signatures in rate-rhythmicity space. Superior colliculus perturbation, for instance, might alter both rate

and rhythmicity simultaneously, unlike the selective rhythmicity effect we observe after FEF inhibition. Transcranial focused ultrasound offers a path toward testing such predictions with the precision needed to target subcortical structures (Deffieux et al., 2013; Farboud et al., 2026; Folloni et al., 2019). Furthermore, the saccade generation network is distributed, and FEF perturbation likely propagates through it; disentangling individual node contributions from network-level compensation remains a challenge for future work. More broadly, because rhythmicity reflects the precision of the generative process rather than its mean output, saccade rhythmicity may prove a sensitive behavioral marker in populations with compromised oculomotor circuits, such as progressive supranuclear palsy or Parkinson's disease.

Acknowledgements

We would like to thank PD Dr. Paul Taylor for helpful discussions on the anatomical localization of the frontal eye field in humans. We would also like to thank Ibrahim Celic, Katrin Knietsch, and Sianna Gösser for help with the data acquisition.



6

Chapter 6

General Discussion

This thesis addressed a shared inferential problem across two otherwise unrelated domains of systems neuroscience: how to recover latent structure from indirect and constrained measurements that do not transparently reveal the process that produced them. The first part focused on the development of machine learning methods for fNIRS data. Chapter 2 introduced Riemannian geometry-based classification for fNIRS and showed that brain states are encoded in the interaction structure between channels. Respecting the inherent geometry of covariance matrices, yielded a substantial gain in classification accuracy over simple traditional univariate methods. Chapter 3 extended this analysis approach to a clinically proof-of-principle paradigm. There, covariance-based classification together with short and individualized motor imagery trials was used to detect command following in healthy controls simulating DoC patients. The second part investigated how latent decomposition can reveal underlying rhythmic saccade timing. Chapter 4 developed a doubly stochastic modelling approach for primate saccade timing, showing that a locally regular rhythmic generator is obscured by slow, non-stationary rate modulation. Decomposing the observed temporal dynamics of saccades to a rate-driven and noise-driven component reconciled previous conflicting findings of rhythmicity in primate saccadic sampling. As a follow-up, Chapter 5 tested this framework causally by inhibiting the right frontal eye field with cTBS. Following inhibition, saccade rhythmicity was increased compared to Vertex control, showing that FEF contributes to saccade timing likely by injecting cycle-to-cycle variability into the local generator. The remainder of the discussion in this thesis takes the two research lines in turn, before briefly drawing out a methodological thread they share. The unity of the two topics proposed is rather inferential than biological: the two lines of work do not converge on a common neuronal mechanism, but on the common recognition that recovering the latent process requires analytical frameworks that are sensitive to its structure rather than to the surface features of the measurement.

Part I: Geometry-aware inference from hemodynamic signals

Chapter 2 showed that machine learning models based on Riemannian geometry outperformed traditional channel-wise feature approaches on fNIRS brain-state decoding. Using Riemannian models nearly doubled eight-way classification accuracy and substantially improved binary classification accuracy across all task pairs (Näher et al., 2025). Chapter 3 showed that the same family of methods transfers to a very different experimental context, with 2 s trials in place of extended recordings, a reduced optode array, binary task-following classification rather than multi-task decoding, and short-separation regression instead of global-component correction. In this clinical proof-of-principle, the Riemannian approach achieved a sensitivity of 100% and a specificity of 89% in detecting command following (Bastian et al., 2025). The two chapters therefore make complementary contributions: Chapter 2 establishes the methodological advantage against appropriate baseline classifiers, and Chapter 3 establishes that the approach retains its utility under the constraints of a clinically motivated design. Both chapters together

provide evidence that expressing multi-channel fNIRS signals using covariance representations and manipulating these representations using their appropriate geometry is a suitable approach to classifying brain states.

Covariance structure as a discriminative representation

The common element across both chapters is that covariance matrices, more so than channel-wise features, served as the input representation for classification. Covariance captures the pairwise statistical dependencies between measurement sites, quantifying how channels relate to each other rather than what each does in isolation. It should be noted though that covariance representations also contain individual channel information, namely on the diagonal of the covariance matrix, because each channel's covariance with itself is simply the channel's variance. The broader concept underlying the choice to include channel interactions in the data representation is a systems neuroscience perspective on brain states: distributed brain states arise from the coordinated activity of many regions and are therefore likely better expressed in the relational structure of that activity across cortical space, rather than in any isolated response per channel (Ebitz & Hayden, 2021; Vyas et al., 2020). Consistent with this view, a long line of distributed-representation work in fMRI and electrophysiology has shown that cognitive and perceptual states are decodable from multivariate activity patterns that no single channel or voxel carries on its own (Haxby et al., 2001; Kriegeskorte et al., 2008; Norman et al., 2006). The two fNIRS chapters presented here demonstrate that the same holds for hemodynamic signals measured optically: covariance carries sufficient state-level information to support accurate classification, thereby extending a representational principle to fNIRS that is already established in EEG-based brain-state inference (Barachant et al., 2012; Congedo et al., 2017).

Chromophore structure shapes the interaction space in fNIRS

Regardless of which method or estimator is used to quantify the interactions between channels, the selection and processing of the time series for each channel that go into this estimation ultimately control what the covariance matrix represents. I will refer to the set of channel interactions derived from a given set of preprocessed time series as the *interaction structure*. In EEG, for example, covariance can be computed on raw broadband time series or on signals filtered into specific frequency bands. Filtering the time series beforehand allows the interaction structure to be defined within a single frequency band that is likely involved in specific neuronal functions. This choice of band selection has direct consequences, because different frequency bands have been linked to distinct neuronal processes. In the primate visual system, for example, feedforward cortico-cortical influences are largely carried by gamma-band (~60–80 Hz) synchronization, whereas feedback influences are carried by alpha/beta-band rhythms (Bastos et al., 2015; Michalareas et al., 2016). Filtering the channel time series to the gamma band therefore restricts the covariance/channel interactions to processes that are well described by that band (e.g. feedforward cortico-cortical coupling). Similarly, filtering the data in alpha/beta would restrict the interaction structure to processes in the alpha-beta band (e.g. feedback).

Covariance can also be computed jointly across gamma and alpha/beta in two different ways. A block-diagonal matrix preserves the within-band interactions for each band separately, while omitting any cross-band terms. A full joint covariance additionally includes cross-frequency entries, which capture coordination between processes operating in different bands, for example between feedforward and feedback signaling. The choice of time series preprocessing and filtering thus determines not only which neural processes contribute to classification, but also whether the interactions between those processes are part of the representation. Constructing an interaction space for a given brain state classification problem is therefore a design step that requires neuroscientific domain knowledge about which processes are expected to differ between the states of interest. Reading this in the other direction, the question of which interaction spaces best discriminate a given pair of states is itself a scientific question, because the answer identifies which neural processes the states differ in. Interaction-space construction is thus both an engineering choice and a tool for neuroscientific inquiry.

fNIRS measures an underlying signal different from EEG and therefore carries different information. Unlike fMRI, which records a single hemodynamic variable, fNIRS records two correlated, but partially independent variables at each measurement site: oxygenated and deoxygenated hemoglobin, whose differential optical absorption is the physical basis of the measurement itself (Ferrari & Quaresima, 2012; Jöbsis, 1977). The corresponding interaction space divides naturally into three families: interactions within HbO, interactions within HbR, and cross-chromophore interactions between HbO and HbR channels. The design choice faced here is directly analogous to the one just described for EEG frequency bands. The question becomes whether to restrict the interaction space to a single chromophore, to combine both chromophores in block-diagonal form that preserves within-chromophore interactions but omits cross-chromophore terms, or to include the cross-chromophore entries in a full joint representation. Chapter 2 established empirically that cross-chromophore interactions contribute little (at least for the set of tasks used in that study) once within-chromophore interactions are already modeled. This is realized through using the block-diagonal form that retains the state-discriminative within-chromophore structure while also allowing the HbO and HbR subspaces to be tuned independently.

Interaction-space design under clinical constraints

Chapter 3 illustrates that the appropriate level of interaction-space control depends on the downstream engineering goal (Bastian et al., 2025). The clinical paradigm in Chapter 3 used binary classification between active imagery and rest. The relevant contrast for classification is the presence or absence of structured task-evoked activity rather than the discrimination of multiple, potentially non-orthogonal, cognitive states. Since the contrast is based on activation vs. rest, any covariance across channel pairs is likely to contribute useful information regardless of whether the pair sits within HbO, within HbR, or across chromophores. Rather than committing to a single interaction space a priori, the clinical tool uses three classifiers that

are trained in parallel on three candidate representations: HbO-only, HbR-only, and a full joint HbO&HbR covariance. The minimum *P*-value across the three corresponding permutation tests determined the diagnostic outcome.

At the group level, the combined HbO&HbR representation produced the highest mean classification accuracy in the responsive condition, consistent with the expectation that a larger covariance dimensionality provides more opportunities to capture structured activation for an activation vs. rest contrast. At the individual level, however, HbO-only or HbR-only carried the strongest evidence in eight of nine participants, with the combined representation winning only for a single participant. Group-level performance and individual-level performance are therefore dissociated, as the representation that wins on average is rarely the representation that wins for any specific participant.

This dissociation has direct implications for the design and deployment of clinical diagnostic methods. Disorders of consciousness arise from heterogeneous brain injuries (Edlow et al., 2021), and the resulting variability in the shape of the hemodynamic response makes it unlikely that any single representation will be optimal across patients. Therefore, it is important that diagnostic approaches in this setting properly accommodate inter-patient variability. This is a fundamentally different design problem from that of an out-of-the-box brain-computer interface that is intended for repeated use by healthy individuals, where consistency across users and sessions is the primary goal. The strategy used in Chapter 3 amounts to running classification on three different interaction spaces in parallel and letting each patient's data select the most informative one. The two fNIRS chapters therefore illustrate complementary uses of the same geometry-aware framework: Chapter 2 shows how to construct a single, chromophore-aware interaction space and tune it carefully, while Chapter 3 shows how to combine evidence across multiple interaction spaces when individual variability makes a single representation unreliable.

Classification as within-session confirmation of task compliance

Chapter 3 also reframes the traditional role of a classifier in this diagnostic context. The successful training of a classifier serves as within-session confirmation that task-related modulation is present, rather than predicting unseen examples. This use of classification as an inference tool is well established in cognitive neuroscience (Hebart & Baker, 2018) and has been the implicit logic of EEG- and fMRI-based command-following paradigms in DoC (Claassen et al., 2019; Cruse et al., 2011; Goldfine et al., 2013). Chapter 3 contributes a specific operationalization of this logic for fNIRS, combining short, individualized trials, covariance-based Riemannian classification, and permutation-based individual-level inference. The logic used for this proof-of-principle generalizes to any paradigm in which the relevant question is whether a participant or patient is complying with a task. Practically, the proposed approach has several advantages such as fast availability of results, modest training data requirements, and, most importantly the tolerance for individual variability in the shape of the hemodynamic response.

The achieved sensitivity and specificity in Chapter 3 should be interpreted in the context of the single false positive result. This misclassification appears to reflect task engagement in a participant that was instructed to remain at rest. This outcome reveals a structural property of the approach: it detects the presence of command-following-related modulation and therefore has an asymmetric failure mode. A participant who engages with the task despite instructions not to comply can produce detectable modulation, whereas a participant who does not engage with the task should not. The clinically relevant question (whether a patient is following commands) is asymmetric in the same way. Positive results from this approach constitute evidence of command-following, whereas negative results are compatible with both non-compliance and undetected compliance.

Limitations

Several limitations apply across both fNIRS studies. The sample sizes were modest in both studies (seven and nine participants, respectively). Both chapters evaluated classifiers using within-participant cross-validation rather than held-out test sets. In the specific applications, cross-validation was appropriate for comparing methods under identical within-participant protocols, and the within-participant design is itself a methodological feature of the clinical proof-of-principle proposed in Chapter 3. Nevertheless, neither of the two studies can provide evidence on how well the proposed approaches generalize across sessions or to unseen participants. Especially in the clinical setting, these insights would be valuable.

Additionally, systemic physiology remains a partial confound in both chapters. fNIRS is sensitive to extracerebral hemodynamic fluctuations arising from scalp blood flow, cardiac pulsation, respiration, and blood pressure changes, all of which can produce task-correlated signals that can be difficult to distinguish from cortical responses (Kirilina et al., 2012; Tachtsidis & Scholkmann, 2016). Chapter 2 lacked short-separation channels and used PCA-based global component regression as a substitute, while Chapter 3 incorporated short-separation channel regression. Neither approach fully excludes the possibility that covariance differences between conditions reflect coordinated extracerebral responses in addition to cortical activity. This caveat constrains the neuronal interpretation of the decoded signal but does not call the decoding performance itself into question.

Per-participant hyperparameter tuning has both positive and negative consequences. It enables clinical individualization and accommodates heterogeneous shapes in the hemodynamic response. Simultaneously, it weakens claims about population-level generalization. Finally, the translational distance from a proof-of-principle simulation of DoC to actual DoC patients is substantial. Therefore, Chapter 3 should be understood as motivating clinical validation rather than providing it.

Outlook

fNIRS has historically been treated as a portable analog of fMRI, with analysis pipelines that were largely inherited from the fMRI tradition (Plichta et al., 2007; Schroeter et al., 2004; Ye et al., 2009). The findings presented here support a shift toward channel interaction-based analysis frameworks that treat channel co-variation as the primary carrier of state information. Future fNIRS studies that use similar classification approaches will likely benefit from high-density recordings that cover a larger cortical surface (Eggebrecht et al., 2014). It also seems plausible to further use hybrid recordings that combine fNIRS with EEG in a natural next step, as both modalities can admit the same covariance-based representation (Ahn & Jun, 2017; Luhmann et al., 2017). Combining optical and electrophysiological measurements within a geometry-aware classification framework would allow the interaction spaces used for state inference to be constructed with a large degree of flexibility that neither modality affords alone. In such a context, the question of which frequency bands, which chromophores, and which cross-modal interactions carry the most discriminative information for a given pair of brain states becomes scientifically interesting in itself, because the answer identifies which neural rhythms, which vascular components, and which cross-modal couplings are recruited for that specific brain state. In this sense, the choice of representation is not a preprocessing step but a commitment about where in the signal the brain-state information is expected to live.

Part II: Two components of primate saccade timing

Chapter 4 showed that primate saccade timing has a compound structure. An intrinsically regular generator operates with a slowly drifting saccade rate. Together, these two components can parsimoniously explain the saccade timing statistics observed during natural viewing. A drift-to-threshold model with only one free noise parameter reproduced the monotonic relationship between saccade rate and local rhythmicity across humans, macaques, and a marmoset (Näher et al., 2023). The decomposition into two processes that affect saccade timing suggests separate origins for each process. The rise-to-threshold dynamics most plausibly reside in brainstem saccade circuitry. Additional top-down modulation that allows for flexible deviations from the baseline rise-to-threshold dynamics could likely come from higher cortical areas. Chapter 5 causally tested this hypothesis. When right FEF was transiently inhibited by continuous theta-burst stimulation (cTBS) (Huang et al., 2005) during free-viewing visual search, saccade rhythmicity increased. Importantly, saccade rate was unaffected. Source-reconstructed resting-state EEG showed three concurrent changes at the site of stimulation: the local $1/f^n$ spectral slope became steeper, oscillatory beta power was attenuated, and alpha-band connectivity between right FEF and the rest of the dorsal attention network was reduced. All behavioral and neuronal effects peaked in the first post-stimulation block and largely returned to pre-intervention levels by the second block. This followed the established temporal signature of cTBS after-effects (Wischnewski & Schutter, 2015). Taken together, the two chapters offer

complementary evidence, but from two different angles: Chapter 4 defines the analytical separation between rate and rhythmicity and proposes a circuit in which a subcortical rhythmic generator is modulated by variable cortical input. Chapter 5 tests this architecture causally and identifies right FEF as one cortical source of the saccade timing variability.

Reconciling prior accounts of saccade timing

Three lines of prior work have produced conclusions that initially look incompatible with the findings of Chapters 4 and 5. However, on closer inspection, each turns out to be compatible once rate fluctuations are accounted for.

The question of whether saccade sequences are rhythmic has been addressed most directly by Amit et al. (2017), who combined spectral peak analysis, hazard function estimation, and ISI shuffling on free-viewing data. They concluded that a self-paced first-order process can sufficiently explain the saccade timing data without the need for a central pacemaker or pattern generator. The simulation results in Chapter 4 revealed that each of these three tools loses power when used on non-stationary processes. Specifically, this means that when ISIs are pooled across time segments with different saccade rates, spectral peaks are smeared by the averaging of different cycle lengths, hazard functions are likewise blurred, and the shuffle test, normally valid under stationarity, does not register the higher-order structure. The Amit et al. (2017) conclusion is therefore correct within its own scope of assuming stationary saccade dynamics. The marginal ISI distribution can be informative, but only once it is recognized as the product of two processes: a local generator and a slowly varying rate. Treating the marginal distribution as a direct window onto the generator collapses this distinction and therefore obscures the rhythmicity that the generator produces at any given rate.

A second mismatch arises with the LATER framework originally formulated by Carpenter and colleagues to account for reaction-time distributions in cued saccade paradigms (Carpenter & Williams, 1995; Noorani & Carpenter, 2016). In this framework, each trial is a linear accumulation toward threshold with a drift rate drawn independently per-trial from a fixed distribution. When the same logic is transferred to free-viewing saccades, successive intervals inherit the same independence assumption. This implies that ISI serial correlations should not extend beyond the refractory period. Chapter 4 showed the opposite, that is, slow drifts in saccade-rate time courses, producing ISI correlations that stretch across multiple saccades. Critically, these long-range correlations disappear once the ISIs are time-rescaled by local rate, indicating that the correlation's origin is the slow saccade rate modulation. The architecture of LATER is preserved at the per-saccade level; what changes is the implementation of drift rate, which becomes a slowly varying control signal rather than an Independent and Identically Distributed (IID) per-cycle variable.

The third mismatch concerns the previous offline TMS literature on FEF, which has generally struggled to produce effects on standard oculomotor summary metrics. Cameron et al. (2015) reported reduced anti-saccade amplitudes and more corrective saccades after right FEF cTBS. However, reaction times and error rates were largely spared. Gurel et al. (2018) applied cTBS to left and right FEF without detecting changes in voluntary or reflexive saccade latencies. Gallotto et al. (2025) targeted the same area and found no cued-attention impairment, a result they attributed to rapid compensation within the attention network. A consistent theme runs through these reports: mean latency, amplitude, and error rate all appear to be preserved when FEF is inhibited by cTBS. This has been read as evidence that remaining oculomotor nodes can compensate for the perturbed node (Cameron et al., 2020). Chapter 5 offers a complementary reading rather than a contradictory one. All those measures describe single saccades, and single-saccade properties could precisely be the aspect which the distributed oculomotor network might be able to compensate. The regularity of a saccade sequence is inherently different, as it reflects cumulative variability across many saccades. To my knowledge, no prior offline-TMS study has examined it. The absence of effects on single-saccade metrics therefore does not imply that FEF inhibition has no behavioral consequence; it implies only that the measures chosen so far do not reveal one.

Sustained cortical inhibition as a tool for sequence-level measurement

Online paradigms, in which pulses are time-locked to saccade preparation, offer trial-by-trial temporal precision: TMS over right FEF delivered during the saccade preparation period prolongs saccade latencies (Nagel et al., 2008; Thickbroom et al., 1996) and paired-pulse TMS over FEF disrupts target discrimination within a narrow interval after stimulus onset (O'Shea et al., 2004). Findings of this kind show that FEF contributes causally to saccade initiation timing. This is consistent with movement-related activity in oculomotor cortex and superior colliculus participating in the rise-to-threshold dynamics that trigger saccades (Dorris et al., 1997; Hanes & Schall, 1996). In the saccade circuit suggested in Chapter 4, this cortical contribution operates on top of a rhythmic subcortical generator: the threshold-reset-refractory cycle in brainstem burst circuitry sets the intrinsic timing of saccade release, and cortical areas such as FEF add variable temporal input that modulates this default cycle (Scudder et al., 2002; Sparks, 2002). Offline cTBS works on a different timescale: one stimulation sequence produces minutes of reduced excitability during which no further pulses are delivered. That feature was indispensable for examining rhythmicity of free-viewing saccade sequences, because TMS pulses delivered during the sequence would themselves impose temporal structure on it and potentially disrupt ongoing rhythmic activity. Online paradigms are well suited for asking when FEF contributes to individual saccades; offline paradigms are required to ask how sustained FEF inhibition reshapes the distribution of intervals that many saccades produce.

A mechanistic picture: subcortical rhythm, cortical variability

Combining previous evidence with the observations reported in Chapters 4 and 5, a more specific circuit for the control of saccade timing can be proposed. The rhythmic component likely arises in subcortical circuitry. Omnipause neurons in the brainstem maintain tonic inhibition of premotor burst neurons during fixation and release it briefly around saccade onset, giving the system a built-in refractory interval after each saccade and implementing a threshold-reset-refractory cycle (Scudder et al., 2002; Sparks, 2002). Movement neurons in both the superior colliculus and FEF converge on a relatively stereotyped discharge level when a saccade is triggered, so that saccade timing depends on how quickly activity climbs toward that threshold (Dorris et al., 1997; Hanes & Schall, 1996). The local regularity that Chapter 4 quantifies follows from this construction, with variability from cycle to cycle determined by noise in the rise to threshold. Cortical projections into this circuit likely act on that noise rather than on other parameters of the cycle. FEF sends dense projections to the superior colliculus and brainstem oculomotor circuits (Segraves & Goldberg, 1987; Stanton et al., 1988) and moment-to-moment changes in FEF activity, driven by task context, stimulus content, or ongoing internal state, would act on individual fixations to match sensory and top-down demands, injecting variability into the drift rate from one cycle to the next. Chapter 5 distinguishes this role from an alternative perspective in which FEF generates or amplifies the rhythm itself. If FEF modulates the rhythm, FEF inhibition should increase rhythmicity; if it is the source of the rhythm, FEF inhibition should decrease it. The observed increase supports the first account. The neuronal changes measured alongside this behavioral effect confirm that the manipulation reduced local FEF excitability and propagated into the wider DAN: a steeper $1/f^n$ slope at the stimulation site and a drop in right-FEF-seeded alpha connectivity.

Saccade rhythmicity and rhythmic attentional sampling

Identifying saccadic eye movements as a rhythmic sampling behavior naturally fits in with the substantial literature on rhythmic covert attention and brain rhythms. Theta-frequency fluctuations in perceptual performance and in the neuronal correlates of attention have been reported in a wide range of different paradigms (Busch et al., 2009; Busch & VanRullen, 2010; Fiebelkorn et al., 2018; Fiebelkorn & Kastner, 2019; Fries, 2023; Landau et al., 2015; Landau & Fries, 2012). Saccades fit in the same temporal regime, since their natural occurrence rate overlaps the frequency bands in which covert-attention paradigms show rhythmic structure, and saccades, or shifts in attention, tend to occur at preferred phases of the attentional rhythm (Fiebelkorn et al., 2018).

An open question in this literature is where the generator is located. Existing evidence has established that frontoparietal regions are involved in the temporal structure of environmental sampling. However, an involvement does not mean that these regions are also the generator of attentional rhythms and correlational evidence alone is insufficient. The present thesis adds a causal point of view. The cTBS effect from Chapter 5, together with the behavioral results from

Chapter 4 suggest that the rhythm's source is likely not in FEF. A subcortical generator under cortical influence, as outlined above, appears to be the more parsimonious explanation.

A functional interpretation: flexibility at the cost of regularity

A mechanistic interpretation of the results in Chapters 4 and 5 suggests that the brainstem and the cortex contribute different aspects to saccade timing. Cortical control likely permits flexibility, which allows fixation durations to track scene content and task demands. The present study shows that this flexibility brings a measurable reduction in temporal regularity/ rhythmicity. Without extensive cortical input, the saccade generation system's default mode is likely highly rhythmic. The appearance of non-rhythmic saccade sequences in healthy primates might therefore reflect the cortical mechanism that interrupts this default.

This architecture has a natural evolutionary logic. Brainstem central pattern generators are evolutionarily old and metabolically cheap: once built, they produce consistent output without continuous supervision, and they suit behaviors that the organism performs repeatedly throughout its life (Grillner, 2006; Marder & Bucher, 2001). Running saccade generation entirely through the cortex would be costly in energy and latency, and potentially fragile. Assigning the default to subcortical circuitry, and adding cortical modulation where needed, gives the system a stable baseline with selective flexibility. The cortex supplies the override that interrupts the default when visual scene or behavioral context calls for it. The capacity of the cortex to integrate sensory context, memory, and behavioral goals makes it the appropriate substrate for such interventions. The same trade-off shapes exploration and exploitation in active sampling: the default rhythm keeps the eyes moving across a stable scene, and cortical modulation slows them down when something is worth dwelling on. The Chapter 5 finding, that inhibiting FEF strengthens saccade rhythmicity, is what this architecture predicts.

Framed this way, overt saccadic sampling takes its place alongside rodent whisking and sniffing, within a larger family of rhythmic active-sensing behaviors (Kleinfeld et al., 2006; Schroeder & Lakatos, 2009; Takatoh et al., 2022). The frontoparietal attention network likely not only contributes to spatial selection in primate saccadic sampling, but to temporal sequencing through the modulation of saccade regularity.

Limitations

Part II of this thesis has several limitations worth discussing. Chapter 4 used a small non-human primate sample (two macaques and one marmoset). Therefore, inferences on the observed cross-species similarity are limited to the sample. The stimulus set also excludes tasks with strong externally imposed temporal structure such as fast and transient movie scenes, under which the rate-rhythmicity decomposition might look different. Chapter 5 used a sample size of sixteen participants. This number was fixed by the Latin-square counterbalancing scheme. cTBS is known to produce heterogeneous responses and has a substantial non-responder rate (Corp

et al., 2020). A larger sample would almost certainly have produced more sensitive estimates. The cTBS inhibition targeted only right FEF, leaving the contribution of left FEF untested. Additionally, the task used in the study was a free-viewing search task, so effects under stronger top-down regimes such as anti-saccade could differ. Finally, the mechanistic inferences in Part II combine generative modeling with causal behavioral perturbation; direct recordings from the brainstem burst generator or from FEF during rate-controlled viewing remain outside the scope of this work.

Outlook

The results from Chapters 4 and 5 suggest several concrete next steps. The drift-to-threshold under cortical influence model predicts different signatures of the rate-rhythmicity coupling, depending on which node of the saccade network is perturbed. Cortical perturbations that reduce drift noise should selectively raise rhythmicity without changing rate (as observed for right FEF here), whereas perturbations of the subcortical generator itself should affect both rate and rhythmicity, because the threshold-crossing machinery is the common substrate for both. Transcranial focused ultrasound, which is non-invasive and can reach subcortical depths (Folloni et al., 2019), is a plausible tool for testing such predictions in the superior colliculus and brainstem saccade circuitry, building on existing FEF work in primates (Deffieux et al., 2013) and humans (Farboud et al., 2026).

The same rate-rhythmicity space may also turn out to be useful as a behavioral marker in populations with compromised oculomotor control. Progressive supranuclear palsy, Parkinson's disease, and other related conditions are usually assessed through mean latency, amplitude, and saccade error rate. The RCRI is orthogonal to these measures and directly captures the integrity of potential cortical modulation of a subcortical generator. Whether this orthogonal dimension is useful in a diagnostic or prognostic setting is an empirical question the present work sets up but does not resolve. Finally, linking behavioral rhythmicity more tightly to its circuit substrate will eventually require direct neuronal recordings from the brainstem burst generator and from other cortical oculomotor areas.

Part III: Latent structure as a unifying inferential move

Despite their methodological and fundamental distance, the two parts of this thesis converge on a common inference problem. In both parts, the saccade and the fNIRS chapters, the main outputs that are being directly measured, are not where the phenomenon of interest lives. The fNIRS field conventionally treats channel-wise amplitudes as the main carrier of brain-state signal. Yet, Chapter 2 shows that the state-discriminative information is more prominent in the interaction structure between channels rather than in any single channel's response. Treating the covariance between channels as the main discriminative feature nearly doubles the eight-

way classification accuracy compared to using channel-wise features. Similarly, saccade research has traditionally treated the marginal ISI distribution as the output of the saccade generator. Yet the rhythmic component of that generator is obscured by a slowly fluctuating saccade rate. The marginal ISI distribution therefore represents an aggregate of many local rate regimes into a single heavy-tailed mixture. Treating the marginal distribution as the direct outcome of the rhythmic generator obscures any local rhythm that the generator produces at a fixed rate.

Different domains require different analysis approaches. Riemannian geometry fits in the fNIRS context, because covariance matrices live on a curved SPD manifold whose geometry should be respected for the manipulation of these matrices. Doubly stochastic gamma-renewal modeling fits for saccade timing models because ISIs share properties of neuronal spike trains, for which rate-rescaling approaches have already been successfully established. Neither analytical approach would transfer directly to the other domain. What transfers however is the recognition that the appropriate inferential approach is the one that matches the latent generative process. Similar inference problems recur across unrelated domains of systems neuroscience, and identifying the right latent structure, whether geometric or generative, is often what resolves the apparent paradox.


A

Appendix

Bibliography
Appendix A: Data Management
Appendix B: Curriculum Vitae
Appendix C: List of Publications
Appendix D: AI Disclosure
Donders Graduate School

Bibliography

Abdalmalak, A., Milej, D., Norton, L., Debicki, D. B., Owen, A. M., & Lawrence, K. S. (2021). The Potential Role of fNIRS in Evaluating Levels of Consciousness. *Frontiers in Human Neuroscience*, *15*. https://doi.org/10.3389/fnhum.2021.703405

Abdalmalak, A., Milej, D., Yip, L. C. M., Khan, A. R., Diop, M., Owen, A. M., & St. Lawrence, K. (2020). Assessing Time-Resolved fNIRS for Brain-Computer Interface Applications of Mental Communication. *Frontiers in Neuroscience*, *14*. https://doi.org/10.3389/fnins.2020.00105

Ahn, S., & Jun, S. C. (2017). Multi-Modal Integration of EEG-fNIRS for Brain-Computer Interfaces – Current Limitations and Future Directions. *Frontiers in Human Neuroscience*, *11*. https://doi.org/10.3389/fnhum.2017.00503

Amit, R., Abeles, D., Bar-Gad, I., & Yuval-Greenberg, S. (2017). Temporal dynamics of saccades explained by a self-paced process. *Scientific Reports*, *7*(1), 886. https://doi.org/10.1038/s41598-017-00881-7

Andrews, K., Murphy, L., Munday, R., & Littlewood, C. (1996). Misdiagnosis of the vegetative state: Retrospective study in a rehabilitation unit. *BMJ*, *313*(7048), 13–16. https://doi.org/10.1136/bmj.313.7048.13

Attwell, D., Buchan, A. M., Charpak, S., Lauritzen, M., MacVicar, B. A., & Newman, E. A. (2010). Glial and neuronal control of brain blood flow. *Nature*, *468*(7321), 232–243. https://doi.org/10.1038/nature09613

Baal, P. van, Morton, A., & Severens, J. L. (2018). Health care input constraints and cost effectiveness analysis decision rules. *Social Science & Medicine (1982)*, *200*, 59. https://doi.org/10.1016/j.socscimed.2018.01.026

Barachant, A., & Bonnet, S. (2011). Channel Selection Procedure using Riemannian distance for BCI applications. *The 5th International IEEE EMBS Conference on Neural Engineering*, TBA. https://hal.science/hal-00602707

Barachant, A., Bonnet, S., Congedo, M., & Jutten, C. (2010). Riemannian geometry applied to BCI classification. In Vigneron, V. ; Zarzoso, V. ; Moreau, E. ; Gribonval, & R. ; Vincent (Eds.), *LVA/ICA 2010—9th International Conference on Latent Variable Analysis and Signal Separation* (Vol. 6365, pp. 629–636). Springer. https://doi.org/10.1007/978-3-642-15995-4_78

Barachant, A., Bonnet, S., Congedo, M., & Jutten, C. (2012). Multiclass Brain–Computer Interface Classification by Riemannian Geometry. *IEEE Transactions on Biomedical Engineering*, *59*(4), 920–928. IEEE Transactions on Biomedical Engineering. https://doi.org/10.1109/TBME.2011.2172210

Barachant, A., Bonnet, S., Congedo, M., & Jutten, C. (2013). Classification of covariance matrices using a Riemannian-based kernel for BCI applications. *Neurocomputing*, *112*, 172–178. https://doi.org/10.1016/j.neucom.2012.12.039

Barachant, A., & Congedo, M. (2014). *A Plug&Play P300 BCI Using Information Geometry* (arXiv:1409.0107). arXiv. https://doi.org/10.48550/arXiv.1409.0107

Bardin, J. C., Fins, J. J., Katz, D. I., Hersh, J., Heier, L. A., Tabelow, K., Dyke, J. P., Ballon, D. J., Schiff, N. D., & Voss, H. U. (2011). Dissociations between behavioural and functional magnetic resonance imaging-based evaluations of cognitive function after brain injury. *Brain*, *134*(3), 769–782. https://doi.org/10.1093/brain/awr005

Bardin, J. C., Schiff, N. D., & Voss, H. U. (2012). Pattern classification of volitional functional magnetic resonance imaging responses in patients with severe brain injury. *Archives of Neurology*, *69*(2), 176–181. https://doi.org/10.1001/archneurol.2011.892

Bastian, L., Näher, T., Vorreuther, A., Lührs, M., Benitez Andonegui, A., Fries, P., Riecke, L., & Sorger, B. (2025). Sensitive and specific fNIRS-based approach for awareness detection in disorders of consciousness: Proof of principle in healthy adults. *Neurophotonics*, *12*(4), 045001. https://doi.org/10.1117/1.NPh.12.4.045001

Bastos, A. M., Vezoli, J., Bosman, C. A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J. R., De Weerd, P., Kennedy, H., & Fries, P. (2015). Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, *85*(2), 390–401. https://doi.org/10.1016/j.neuron.2014.12.018

Bellet, J., Chen, C.-Y., & Hafed, Z. M. (2017). Sequential hemifield gating of α- and β-behavioral performance oscillations after microsaccades. *Journal of Neurophysiology*, *118*(5), 2789–2805. https://doi.org/10.1152/jn.00253.2017

Bellet, M. E., Bellet, J., Nienborg, H., Hafed, Z. M., & Berens, P. (2019). Human-level saccade detection performance using deep neural networks. *Journal of Neurophysiology*, *121*(2), 646–661. https://doi.org/10.1152/jn.00601.2018

Benedetto, A., & Morrone, M. C. (2017). Saccadic Suppression Is Embedded Within Extended Oscillatory Modulation of Sensitivity. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *37*(13), 3661–3670. https://doi.org/10.1523/JNEUROSCI.2390-16.2016

Benitez-Andonegui, A., Lührs, M., Nagels-Coune, L., Ivanov, D., Goebel, R., & Sorger, B. (2021). Guiding functional near-infrared spectroscopy optode-layout design using individual (f)MRI data: Effects on signal strength. *Neurophotonics*, *8*(2), 025012. https://doi.org/10.1117/1.NPh.8.2.025012

Biswal, B., Yetkin, F. Z., Haughton, V. M., & Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance in Medicine*, *34*(4), 537–541. https://doi.org/10.1002/mrm.1910340409

Bompas, A., & Sumner, P. (2011). Saccadic Inhibition Reveals the Timing of Automatic and Voluntary Signals in the Human Brain. *Journal of Neuroscience*, *31*(35), 12501–12512. https://doi.org/10.1523/JNEUROSCI.2234-11.2011

Bosman, C. A., Schoffelen, J.-M., Brunet, N., Oostenveld, R., Bastos, A. M., Womelsdorf, T., Rubehn, B., Stieglitz, T., De Weerd, P., & Fries, P. (2012). Attentional Stimulus Selection through Selective Synchronization between Monkey Visual Areas. *Neuron*, *75*(5), 875–888. https://doi.org/10.1016/j.neuron.2012.06.037

Bray, T. J. P., & Carpenter, R. H. S. (2015). Saccadic foraging: Reduced reaction time to informative targets. *The European Journal of Neuroscience*, *41*(7), 908–913. https://doi.org/10.1111/ejn.12845

Breen, E. J., & Williams, K. L. (1990). Contrast matching in a sequence of images by histogram normalization. *Journal of Microscopy*, *159*(2), 223–227. https://doi.org/10.1111/j.1365-2818.1990.tb04779.x

Buonocore, A., & Hafed, Z. M. (2023). The inevitability of visual interruption. *Journal of Neurophysiology*, *130*(2), 225–237. https://doi.org/10.1152/jn.00441.2022

Busch, N. A., Dubois, J., & VanRullen, R. (2009). The Phase of Ongoing EEG Oscillations Predicts Visual Perception. *Journal of Neuroscience*, *29*(24), 7869–7876. https://doi.org/10.1523/JNEUROSCI.0113-09.2009

Busch, N. A., & VanRullen, R. (2010). Spontaneous EEG oscillations reveal periodic sampling of visual attention. *Proceedings of the National Academy of Sciences*, *107*(37), 16048–16053. https://doi.org/10.1073/pnas.1004801107

Cameron, I. G. M., Cretu, A. L., Struik, F., & Toni, I. (2020). The Effects of a TMS Double Perturbation to a Cortical Network. *eNeuro*, *7*(1). https://doi.org/10.1523/ENEURO.0188-19.2019

Cameron, I. G. M., Riddle, J. M., & D'Esposito, M. (2015). Dissociable Roles of Dorsolateral Prefrontal Cortex and Frontal Eye Fields During Saccadic Eye Movements. *Frontiers in Human Neuroscience*, *9*. https://doi.org/10.3389/fnhum.2015.00613

Capotosto, P., Babiloni, C., Romani, G. L., & Corbetta, M. (2009). Frontoparietal Cortex Controls Spatial Attention through Modulation of Anticipatory Alpha Rhythms. *Journal of Neuroscience*, *29*(18), 5863–5872. https://doi.org/10.1523/JNEUROSCI.0539-09.2009

Carpenter, R. H. S., & Williams, M. L. L. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature*, *377*(6544), 59–62. https://doi.org/10.1038/377059a0

Cazzoli, D., Jung, S., Nyffeler, T., Nef, T., Wurtz, P., Mosimann, U. P., & Müri, R. M. (2015). The role of the right frontal eye field in overt visual attention deployment as assessed by free visual exploration. *Neuropsychologia*, *74*, 37–41. https://doi.org/10.1016/j.neuropsychologia.2015.01.027

Chen, Y., Wiesel, A., Eldar, Y. C., & III, A. O. H. (2009). *Shrinkage Algorithms for MMSE Covariance Estimation* (arXiv:0907.4698). arXiv. https://doi.org/10.48550/arXiv.0907.4698

Chevallier, S., Kalunga, E. K., Barthélemy, Q., & Monacelli, E. (2021). Review of Riemannian Distances and Divergences, Applied to SSVEP-based BCI. *Neuroinformatics*, *19*(1), 93–106. https://doi.org/10.1007/s12021-020-09473-9

Childs, N. L., Mercer, W. N., & Childs, H. W. (1993). Accuracy of diagnosis of persistent vegetative state. *Neurology*, *43*(8), 1465–1465. https://doi.org/10.1212/WNL.43.8.1465

Claassen, J., Doyle, K., Matory, A., Couch, C., Burger, K. M., Velazquez, A., Okonkwo, J. U., King, J.-R., Park, S., Agarwal, S., Roh, D., Megjhani, M., Eliseyev, A., Connolly, E. S., & Rohaut, B. (2019). Detection of Brain Activation in Unresponsive Patients with Acute Brain Injury. *The New England Journal of Medicine*, *380*(26), 2497–2505. https://doi.org/10.1056/NEJMoa1812757

Cohen, M. S., Kosslyn, S. M., Breiter, H. C., DiGirolamo, G. J., Thompson, W. L., Anderson, A. K., Bookheimer, S. Y., Rosen, B. R., & Belliveau, J. W. (1996). Changes in cortical activity during mental rotation A mapping study using functional MRI. *Brain*, *119*(1), 89–100. https://doi.org/10.1093/brain/119.1.89

Colclough, G. L., Brookes, M. J., Smith, S. M., & Woolrich, M. W. (2015). A symmetric multivariate leakage correction for MEG connectomes. *NeuroImage*, *117*, 439–448. https://doi.org/10.1016/j.neuroimage.2015.03.071

Coleman, M. R., Davis, M. H., Rodd, J. M., Robson, T., Ali, A., Owen, A. M., & Pickard, J. D. (2009). Towards the routine use of brain imaging to aid the clinical diagnosis of disorders of consciousness. *Brain: A Journal of Neurology*, *132*(Pt 9), 2541–2552. https://doi.org/10.1093/brain/awp183

Combrisson, E., & Jerbi, K. (2015). Exceeding chance level by chance: The caveat of theoretical chance levels in brain signal classification and statistical assessment of decoding accuracy. *Journal of Neuroscience Methods*, *250*, 126–136. https://doi.org/10.1016/j.jneumeth.2015.01.010

Congedo, M., Barachant, A., & Bhatia, R. (2017). Riemannian geometry for EEG-based brain-computer interfaces; a primer and a review. *Brain-Computer Interfaces*, *4*(3), 155–174. https://doi.org/10.1080/2326263X.2017.1297192

Corp, D. T., Bereznicki, H. G. K., Clark, G. M., Youssef, G. J., Fried, P. J., Jannati, A., Davies, C. B., Gomes-Osman, J., Stamm, J., Chung, S. W., Bowe, S. J., Rogasch, N. C., Fitzgerald, P. B., Koch, G., Lazzaro, V. D., Pascual-Leone, A., & Enticott, P. G. (2020). Large-scale analysis of interindividual variability in theta-burst stimulation data:

Results from the 'Big TMS Data Collaboration.' *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *13*(5), 1476–1488. https://doi.org/10.1016/j.brs.2020.07.018

Cruse, D., Chennu, S., Chatelle, C., Bekinschtein, T. A., Fernández-Espejo, D., Pickard, J. D., Laureys, S., & Owen, A. M. (2011). Bedside detection of awareness in the vegetative state: A cohort study. *The Lancet*, *378*(9809), 2088–2094. https://doi.org/https://doi.org/10.1016/S0140-6736(11)61224-5

Deffieux, T., Younan, Y., Wattiez, N., Tanter, M., Pouget, P., & Aubry, J.-F. (2013). Low-intensity focused ultrasound modulates monkey visuomotor behavior. *Current Biology: CB*, *23*(23), 2430–2433. https://doi.org/10.1016/j.cub.2013.10.029

Delgado, J. M. C., Achanccaray, D., Villota, E. R., & Chevallier, S. (2020). Riemann-Based Algorithms Assessment for Single- and Multiple-Trial P300 Classification in Non-Optimal Environments. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *28*(12), 2754–2761. IEEE Transactions on Neural Systems and Rehabilitation Engineering. https://doi.org/10.1109/TNSRE.2020.3043418

Dorris, M. C., Paré, M., & Munoz, D. P. (1997). Neuronal Activity in Monkey Superior Colliculus Related to the Initiation of Saccadic Eye Movements. *Journal of Neuroscience*, *17*(21), 8566–8579. https://doi.org/10.1523/JNEUROSCI.17-21-08566.1997

Drakaki, M., Mathiesen, C., Siebner, H. R., Madsen, K., & Thielscher, A. (2022). Database of 25 validated coil models for electric field simulations for TMS. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation*, *15*(3), 697–706. https://doi.org/10.1016/j.brs.2022.04.017

Duecker, F., Formisano, E., & Sack, A. T. (2013). Hemispheric Differences in the Voluntary Control of Spatial Attention: Direct Evidence for a Right-Hemispheric Dominance within Frontal Cortex. *Journal of Cognitive Neuroscience*, *25*(8), 1332–1342. https://doi.org/10.1162/jocn_a_00402

Duecker, F., & Sack, A. T. (2015). The hybrid model of attentional control: New insights into hemispheric asymmetries inferred from TMS research. *Neuropsychologia, Effects of Non-Invasive Brain Stimulation on Attention: Current Debates, Cognitive Studies and Novel Clinical Applications*, *74*, 21–29. https://doi.org/10.1016/j.neuropsychologia.2014.11.023

Dugué, L., Beck, A.-A., Marque, P., & VanRullen, R. (2019). Contribution of FEF to Attentional Periodicity during Visual Search: A TMS Study. *eNeuro*, *6*(3). https://doi.org/10.1523/ENEURO.0357-18.2019

Eastmond, C., Subedi, A., De, S., & Intes, X. (2022). Deep learning in fNIRS: A review. *Neurophotonics*, *9*(04). https://doi.org/10.1117/1.NPh.9.4.041411

Ebitz, R. B., & Hayden, B. Y. (2021). The population doctrine in cognitive neuroscience. *Neuron*, *109*(19), 3055–3068. https://doi.org/10.1016/j.neuron.2021.07.011

Edlow, B. L., Claassen, J., Schiff, N. D., & Greer, D. M. (2021). Recovery from disorders of consciousness: Mechanisms, prognosis and emerging therapies. *Nature Reviews Neurology*, *17*(3), 135–156. https://doi.org/10.1038/s41582-020-00428-x

Eggebrecht, A. T., Ferradal, S. L., Robichaux-Viehoever, A., Hassanpour, M. S., Dehghani, H., Snyder, A. Z., Hershey, T., & Culver, J. P. (2014). Mapping distributed brain function and networks with diffuse optical tomography. *Nature Photonics*, *8*(6), 448–454. https://doi.org/10.1038/nphoton.2014.107

Engbert, R., & Kliegl, R. (2003). Microsaccades uncover the orientation of covert attention. *Vision Research*, *43*(9), 1035–1045. https://doi.org/10.1016/S0042-6989(03)00084-1

Epstein, R. A., Patai, E. Z., Julian, J. B., & Spiers, H. J. (2017). The cognitive map in humans: Spatial navigation and beyond. *Nature Neuroscience*, *20*(11), 1504–1513. https://doi.org/10.1038/nn.4656

Fabbri, F., Sassaroli, A., Henry, M. E., & Fantini, S. (2004). Optical measurements of absorption changes in two-layered diffusive media. *Physics in Medicine and Biology*, *49*(7), 1183–1201. https://doi.org/10.1088/0031-9155/49/7/007

Farboud, S., Kop, B. R., Koolschijn, R. S., Walstra, S. L. Y., Marques, J. P., Chetverikov, A., Medendorp, W. P., Verhagen, L., & den Ouden, H. E. M. (2026). Rapid modulation of choice behavior by ultrasound on the human frontal eye fields. *Nature Communications*, *17*(1), 2966. https://doi.org/10.1038/s41467-026-69854-7

Fernández-Espejo, D., Norton, L., & Owen, A. M. (2014). The Clinical Utility of fMRI for Identifying Covert Awareness in the Vegetative State: A Comparison of Sensitivity between 3T and 1.5T. *PLOS ONE*, *9*(4), e95082. https://doi.org/10.1371/journal.pone.0095082

Ferrari, M., & Quaresima, V. (2012). A brief review on the history of human functional near-infrared spectroscopy (fNIRS) development and fields of application. *NeuroImage*, *63*(2), 921–935. https://doi.org/10.1016/j.neuroimage.2012.03.049

Fiebelkorn, I. C., & Kastner, S. (2019). A Rhythmic Theory of Attention. *Trends in Cognitive Sciences*, *23*(2), 87–101. https://doi.org/10.1016/j.tics.2018.11.009

Fiebelkorn, I. C., Pinsk, M. A., & Kastner, S. (2018). A Dynamic Interplay within the Frontoparietal Network Underlies Rhythmic Spatial Attention. *Neuron*, *99*(4), 842-853.e8. https://doi.org/10.1016/j.neuron.2018.07.038

Fiebelkorn, I. C., Saalmann, Y. B., & Kastner, S. (2013). Rhythmic Sampling within and between Objects despite Sustained Attention at a Cued Location. *Current Biology*, *23*(24), 2553–2558. https://doi.org/10.1016/j.cub.2013.10.063

Fishburn, F. A., Ludlum, R. S., Vaidya, C. J., & Medvedev, A. V. (2019). Temporal Derivative Distribution Repair (TDDR): A motion correction method for fNIRS. *NeuroImage*, *184*, 171–179. https://doi.org/10.1016/j.neuroimage.2018.09.025

Fletcher, P. T., Lu, C., Pizer, S. M., & Joshi, S. (2004). Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging*, *23*(8), 995–1005. IEEE Transactions on Medical Imaging. https://doi.org/10.1109/TMI.2004.831793

Folloni, D., Verhagen, L., Mars, R. B., Fouragnan, E., Constans, C., Aubry, J.-F., Rushworth, M. F. S., & Sallet, J. (2019). Manipulation of Subcortical and Deep Cortical Activity in the Primate Brain Using Transcranial Focused Ultrasound Stimulation. *Neuron*, *101*(6), 1109-1116.e5. https://doi.org/10.1016/j.neuron.2019.01.019

Formont, P., Ovarlez, J.-P., & Pascal, F. (2013). On the use of Matrix Information Geometry for Polarimetric SAR Image Classification. In *Matrix Information Geometry* (pp. 257–276). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-30232-9_10

Fox, M. D., & Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, *8*(9), 700–711. https://doi.org/10.1038/nrn2201

Freedman, S., Golberstein, E., Huang, T.-Y., Satin, D. J., & Smith, L. B. (2021). Docs with their eyes on the clock? The effect of time pressures on primary care productivity. *Journal of Health Economics*, *77*, 102442. https://doi.org/10.1016/j.jhealeco.2021.102442

Fries, P. (2023). Rhythmic attentional scanning. *Neuron*, *111*(7), 954–970. https://doi.org/10.1016/j.neuron.2023.02.015

Fries, P., & Maris, E. (2022). What to Do If N Is Two? *Journal of Cognitive Neuroscience*, *34*(7), 1114–1118. https://doi.org/10.1162/jocn_a_01857

Fries, P., Womelsdorf, T., Oostenveld, R., & Desimone, R. (2008). The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *28*(18), 4823–4835. https://doi.org/10.1523/JNEUROSCI.4499-07.2008

Galiotta, V., Quattrociocchi, I., D'Ippolito, M., Schettini, F., Aricò, P., Sdoia, S., Formisano, R., Cincotti, F., Mattia, D., & Riccio, A. (2022). EEG-based Brain-Computer Interfaces for people with Disorders of Consciousness: Features and applications. A systematic review. *Frontiers in Human Neuroscience*, *16*, 1040816. https://doi.org/10.3389/fnhum.2022.1040816

Gallegos-Ayala, G., Furdea, A., Takano, K., Ruf, C. A., Flor, H., & Birbaumer, N. (2014). Brain communication in a completely locked-in patient using bedside near-infrared spectroscopy. *Neurology*, *82*(21), 1930–1932. https://doi.org/10.1212/WNL.0000000000000449

Gallotto, S., Schuhmann, T., Duecker, F., & Sack, A. T. (2025). Top-down attentional control resilience reveals time-sensitive transcranial magnetic stimulation effects. *Cerebral Cortex*, *35*(7), bhaf168. https://doi.org/10.1093/cercor/bhaf168

Gardumi, A., Ivanov, D., Hausfeld, L., Valente, G., Formisano, E., & Uludağ, K. (2016). The effect of spatial resolution on decoding accuracy in fMRI multivariate pattern analysis. *NeuroImage*, *132*, 32–42. https://doi.org/10.1016/j.neuroimage.2016.02.033

Gaur, P., Pachori, R. B., Wang, H., & Prasad, G. (2018). A Multi-class EEG-based BCI classification using Multivariate Empirical Mode Decomposition Based Filtering and Riemannian Geometry. *Expert Systems with Applications*, *95*, 201–211. https://doi.org/10.1016/j.eswa.2017.11.007

Gemignani, J., Middell, E., Barbour, R. L., Graber, H. L., & Blankertz, B. (2018). Improving the analysis of near-infrared spectroscopy data with multivariate classification of hemodynamic patterns: A theoretical formulation and validation. *Journal of Neural Engineering*, *15*(4), 045001. https://doi.org/10.1088/1741-2552/aabb7c

Giacino, J. T., Kalmar, K., & Whyte, J. (2004). The JFK Coma Recovery Scale-Revised: Measurement characteristics and diagnostic utility. *Archives of Physical Medicine and Rehabilitation*, *85*(12), 2020–2029. https://doi.org/10.1016/j.apmr.2004.02.033

Gibson, R. M., Fernández-Espejo, D., Gonzalez-Lara, L. E., Kwan, B. Y., Lee, D. H., Owen, A. M., & Cruse, D. (2014). Multiple tasks and neuroimaging modalities increase the likelihood of detecting covert awareness in patients with disorders of consciousness. *Frontiers in Human Neuroscience*, *8*, 950. https://doi.org/10.3389/fnhum.2014.00950

Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C. F., Jenkinson, M., Smith, S. M., & Van Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, *536*(7615), 171–178. https://doi.org/10.1038/nature18933

Goldfine, A. M., Bardin, J. C., Noirhomme, Q., Fins, J. J., Schiff, N. D., & Victor, J. D. (2013). Reanalysis of "Bedside detection of awareness in the vegetative state: A cohort study." *The Lancet*, *381*(9863), 289–291. https://doi.org/10.1016/S0140-6736(13)60125-7

Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., Goj, R., Jas, M., Brooks, T., Parkkonen, L., & Hämäläinen, M. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, *7*, 267. https://doi.org/10.3389/fnins.2013.00267

Grillner, S. (2006). Biological pattern generation: The cellular and computational logic of networks in motion. *Neuron*, *52*(5), 751–766. https://doi.org/10.1016/j.neuron.2006.11.008

Gurel, S. C., Castelo-Branco, M., Sack, A. T., & Duecker, F. (2018). Assessing the Functional Role of Frontal Eye Fields in Voluntary and Reflexive Saccades Using Continuous Theta Burst Stimulation. *Frontiers in Neuroscience*, *12*. https://doi.org/10.3389/fnins.2018.00944

Haeussinger, F. B., Dresler, T., Heinzel, S., Schecklmann, M., Fallgatter, A. J., & Ehlis, A.-C. (2014). Reconstructing functional near-infrared spectroscopy (fNIRS) signals impaired by extra-cranial confounds: An easy-to-use filter method. *NeuroImage*, *95*, 69–79. https://doi.org/10.1016/j.neuroimage.2014.02.035

Hafed, Z. M. (2013). Alteration of Visual Perception prior to Microsaccades. *Neuron*, *77*(4), 775–786. https://doi.org/10.1016/j.neuron.2012.12.014

Hafed, Z. M., & Clark, J. J. (2002). Microsaccades as an overt measure of covert attention shifts. *Vision Research*, *42*(22), 2533–2545. https://doi.org/10.1016/S0042-6989(02)00263-8

Hafed, Z. M., & Ignashchenkova, A. (2013). On the Dissociation between Microsaccade Rate and Direction after Peripheral Cues: Microsaccadic Inhibition Revisited. *The Journal of Neuroscience*, *33*(41), 16220–16235. https://doi.org/10.1523/JNEUROSCI.2240-13.2013

Hallett, P. E., & Adams, B. D. (1980). The predictability of saccadic latency in a novel voluntary oculomotor task. *Vision Research*, *20*(4), 329–339. https://doi.org/10.1016/0042-6989(80)90019-X

Hallett, P. E., & Lightstone, A. D. (1976). Saccadic eye movements towards stimuli triggered by prior saccades. *Vision Research*, *16*(1), 99–106. https://doi.org/10.1016/0042-6989(76)90083-3

Hamada, M., Murase, N., Hasan, A., Balaratnam, M., & Rothwell, J. C. (2013). The Role of Interneuron Networks in Driving Human Motor Cortical Plasticity. *Cerebral Cortex*, *23*(7), 1593–1605. https://doi.org/10.1093/cercor/bhs147

Hanes, D. P., & Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science*, *274*(5286), 427–430. https://doi.org/10.1126/science.274.5286.427

Harandi, M., Salzmann, M., & Hartley, R. (2018). Dimensionality Reduction on SPD Manifolds: The Emergence of Geometry-Aware Methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *40*(1), 48–62. https://doi.org/10.1109/TPAMI.2017.2655048

Hassan, U., Pillen, S., Zrenner, C., & Bergmann, T. O. (2022). The Brain Electrophysiological recording & STimulation (BEST) toolbox. *Brain Stimulation*, *15*(1), 109–115. https://doi.org/10.1016/j.brs.2021.11.017

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The Elements of Statistical Learning*. Springer. https://doi.org/10.1007/978-0-387-84858-7

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, *293*(5539), 2425–2430. https://doi.org/10.1126/science.1063736

Hebart, M. N., & Baker, C. I. (2018). Deconstructing multivariate decoding for the study of brain function. *NeuroImage, New Advances in Encoding and Decoding of Brain Signals*, *180*, 4–18. https://doi.org/10.1016/j.neuroimage.2017.08.005

Herff, C., Heger, D., Putze, F., Hennrich, J., Fortmann, O., & Schultz, T. (2013). Classification of mental tasks in the prefrontal cortex using fNIRS. *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2160–2163. https://doi.org/10.1109/EMBC.2013.6609962

Hipp, J. F., Hawellek, D. J., Corbetta, M., Siegel, M., & Engel, A. K. (2012). Large-scale cortical correlation structure of spontaneous oscillatory activity. *Nature Neuroscience*, *15*(6), 884–890. https://doi.org/10.1038/nn.3101

Holper, L., Biallas, M., & Wolf, M. (2009). Task complexity relates to activation of cortical motor areas during uni- and bimanual performance: A functional NIRS study. *NeuroImage*, *46*(4), 1105–1113. https://doi.org/10.1016/j.neuroimage.2009.03.027

Holt, G. R., Softky, W. R., Koch, C., & Douglas, R. J. (1996). Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. *Journal of Neurophysiology*, *75*(5), 1806–1814. https://doi.org/10.1152/jn.1996.75.5.1806

Hong, K.-S., Khan, M. J., & Hong, M. J. (2018). Feature Extraction and Classification Methods for Hybrid fNIRS-EEG Brain-Computer Interfaces. *Frontiers in Human Neuroscience*, *12*. https://doi.org/10.3389/fnhum.2018.00246

Huang, Y.-Z., Edwards, M. J., Rounis, E., Bhatia, K. P., & Rothwell, J. C. (2005). Theta burst stimulation of the human motor cortex. *Neuron*, *45*(2), 201–206. https://doi.org/10.1016/j.neuron.2004.12.033

Hung, J., Driver, J., & Walsh, V. (2011). Visual Selection and the Human Frontal Eye Fields: Effects of Frontal Transcranial Magnetic Stimulation on Partial Report Analyzed by Bundesen's Theory of Visual Attention. *Journal of Neuroscience*, *31*(44), 15904–15913. https://doi.org/10.1523/JNEUROSCI.2626-11.2011

Iadecola, C. (2017). The Neurovascular Unit Coming of Age: A Journey through Neurovascular Coupling in Health and Disease. *Neuron*, *96*(1), 17–42. https://doi.org/10.1016/j.neuron.2017.07.030

Irimia, A., & Horn, J. D. V. (2015). Functional neuroimaging of traumatic brain injury: Advances and clinical utility. *Neuropsychiatric Disease and Treatment*, *11*, 2355. https://doi.org/10.2147/NDT.S79174

Izzuddin, T. A., Safri, N. M., & Othman, M. A. (2021). Mental imagery classification using one-dimensional convolutional neural network for target selection in single-channel BCI-controlled mobile robot. *Neural Computing and Applications*, *33*(11), 6233–6246. https://doi.org/10.1007/s00521-020-05393-6

Jöbsis, F. F. (1977). Noninvasive, Infrared Monitoring of Cerebral and Myocardial Oxygen Sufficiency and Circulatory Parameters. *Science*, *198*(4323), 1264–1267. https://doi.org/10.1126/science.929199

Kazazian, K., Abdalmalak, A., Novi, S. L., Norton, L., Moulavi-Ardakani, R., Kolisnyk, M., Gofton, T. E., Mesquita, R. C., Owen, A. M., & Debicki, D. B. (2024). Functional near-infrared spectroscopy: A novel tool for detecting consciousness after acute severe brain injury. *Proceedings of the National Academy of Sciences*, *121*(36), e2402723121. https://doi.org/10.1073/pnas.2402723121

Kempny, A. M., James, L., Yelden, K., Duport, S., Farmer, S., Playford, E. D., & Leff, A. P. (2016). Functional near infrared spectroscopy as a probe of brain function in people with prolonged disorders of consciousness. *NeuroImage: Clinical*, *12*, 312–319. https://doi.org/10.1016/j.nicl.2016.07.013

Keshmiri, S., Sumioka, H., Okubo, M., & Ishiguro, H. (2019). An Information-Theoretic Approach to Quantitative Analysis of the Correspondence Between Skin Blood Flow and Functional Near-Infrared Spectroscopy Measurement in Prefrontal Cortex Activity. *Frontiers in Neuroscience*, *13*. https://doi.org/10.3389/fnins.2019.00079

Kirilina, E., Jelzow, A., Heine, A., Niessing, M., Wabnitz, H., Brühl, R., Ittermann, B., Jacobs, A. M., & Tachtsidis, I. (2012). The physiological origin of task-evoked systemic artefacts in functional near infrared spectroscopy. *NeuroImage*, *61*(1), 70–81. https://doi.org/10.1016/j.neuroimage.2012.02.074

Klein, F., Kohl, S. H., Lührs, M., Mehler, D. M. A., & Sorger, B. (2024). From lab to life: Challenges and perspectives of fNIRS for haemodynamic-based neurofeedback in real-world environments. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *379*(1915), 20230087. https://doi.org/10.1098/rstb.2023.0087

Kleinfeld, D., Ahissar, E., & Diamond, M. E. (2006). Active sensation: Insights from the rodent vibrissa sensorimotor system. *Current Opinion in Neurobiology*, *16*(4), 435–444. https://doi.org/10.1016/j.conb.2006.06.009

Koch, S. P., Werner, P., Steinbrink, J., Fries, P., & Obrig, H. (2009). Stimulus-Induced and State-Dependent Sustained Gamma Activity Is Tightly Coupled to the Hemodynamic Response in Humans. *Journal of Neuroscience*, *29*(44), 13962–13970. https://doi.org/10.1523/JNEUROSCI.1402-09.2009

Kohl, S. H., Mehler, D. M. A., Lührs, M., Thibault, R. T., Konrad, K., & Sorger, B. (2020). The Potential of Functional Near-Infrared Spectroscopy-Based Neurofeedback—A Systematic Review and Recommendations for Best Practice. *Frontiers in Neuroscience*, *14*, 594. https://doi.org/10.3389/fnins.2020.00594

Kohno, S., Miyai, I., Seiyama, A., Oda, I., Ishikawa, A., Tsuneishi, S., Amita, T., & Shimizu, K. (2007). Removal of the skin blood flow artifact in functional near-infrared spectroscopic imaging data through independent component analysis. *Journal of Biomedical Optics*, *12*(6), 062111. https://doi.org/10.1117/1.2814249

Kondziella, D., Bender, A., Diserens, K., van Erp, W., Estraneo, A., Formisano, R., Laureys, S., Naccache, L., Ozturk, S., Rohaut, B., Sitt, J. D., Stender, J., Tiainen, M., Rossetti, A. O., Gosseries, O., Chatelle, C., & EAN Panel on Coma, Disorders of Consciousness. (2020). European Academy of Neurology guideline on the diagnosis of coma and other disorders of consciousness. *European Journal of Neurology*, *27*(5), 741–756. https://doi.org/10.1111/ene.14151

Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis—Connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, *2*. https://doi.org/10.3389/neuro.06.004.2008

Kumar, S., Mamun, K., & Sharma, A. (2017). CSP-TSM: Optimizing the performance of Riemannian tangent space mapping using common spatial pattern for MI-BCI. *Computers in Biology and Medicine*, *91*, 231–242. https://doi.org/10.1016/j.compbiomed.2017.10.025

Kurz, E.-M., Wood, G., Kober, S. E., Schippinger, W., Pichler, G., Müller-Putz, G., & Bauernfeind, G. (2018). Towards using fNIRS recordings of mental arithmetic for the detection of residual cognitive activity in patients with disorders of consciousness (DOC). *Brain and Cognition*, *125*, 78–87. https://doi.org/10.1016/j.bandc.2018.06.002

Landau, A. N., & Fries, P. (2012). Attention Samples Stimuli Rhythmically. *Current Biology*, *22*(11), 1000–1004. https://doi.org/10.1016/j.cub.2012.03.054

Landau, A. N., Schreyer, H. M., van Pelt, S., & Fries, P. (2015). Distributed Attention Is Implemented through Theta-Rhythmic Gamma Modulation. *Current Biology*, *25*(17), 2332–2337. https://doi.org/10.1016/j.cub.2015.07.048

Laureys, S., Celesia, G. G., Cohadon, F., Lavrijsen, J., León-Carrión, J., Sannita, W. G., Sazbon, L., Schmutzhard, E., von Wild, K. R., Zeman, A., Dolce, G., & the European Task Force on Disorders of Consciousness. (2010).

Unresponsive wakefulness syndrome: A new name for the vegetative state or apallic syndrome. *BMC Medicine*, *8*(1), 68. https://doi.org/10.1186/1741-7015-8-68

Li, M., Yang, Y., Zhang, Y., Gao, Y., Jing, R., Dang, Y., Chen, X., He, J., & Si, J. (2021). Detecting Residual Awareness in Patients With Prolonged Disorders of Consciousness: An fNIRS Study. *Frontiers in Neurology*, *12*, 618055. https://doi.org/10.3389/fneur.2021.618055

Liu, Y., Kang, X.-G., Chen, B.-B., Song, C.-G., Liu, Y., Hao, J.-M., Yuan, F., & Jiang, W. (2023). Detecting residual brain networks in disorders of consciousness: A resting-state fNIRS study. *Brain Research*, *1798*, 148162. https://doi.org/10.1016/j.brainres.2022.148162

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, *453*(7197), 869–878. https://doi.org/10.1038/nature06976

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, *412*(6843), 150–157. https://doi.org/10.1038/35084005

Lu, H., Jiang, J., Si, J., Wang, Y., & Huang, F. (2023). A functional near-infrared spectroscopy study on hemodynamic changes of patients with prolonged disorders of consciousness responding to different auditory stimuli. *BMC Neurology*, *23*(1), 242. https://doi.org/10.1186/s12883-023-03292-6

Luauté, J., Morlet, D., & Mattout, J. (2015). BCI in patients with disorders of consciousness: Clinical perspectives. *Annals of Physical and Rehabilitation Medicine, Brain Computer Interfaces (BCIs) / Coordinated by Jacques Luauté and Isabelle Laffont*, *58*(1), 29–34. https://doi.org/10.1016/j.rehab.2014.09.015

Luhmann, A., Wabnitz, H., Sander, T., & Müller, K.-R. (2017). M3BA: A Mobile, Modular, Multimodal Biosignal Acquisition Architecture for Miniaturized EEG-NIRS-Based Hybrid BCI and Monitoring. *IEEE Transactions on Bio-Medical Engineering*, *64*(6), 1199–1210. https://doi.org/10.1109/TBME.2016.2594127

Maher, A., Mian Qaisar, S., Salankar, N., Jiang, F., Tadeusiewicz, R., Pławiak, P., Abd El-Latif, A. A., & Hammad, M. (2023). Hybrid EEG-fNIRS brain-computer interface based on the non-linear features extraction and stacking ensemble learning. *Biocybernetics and Biomedical Engineering*, *43*(2), 463–475. https://doi.org/10.1016/j.bbe.2023.05.001

Manning, J. R., Jacobs, J., Fried, I., & Kahana, M. J. (2009). Broadband Shifts in Local Field Potential Power Spectra Are Correlated with Single-Neuron Spiking in Humans. *Journal of Neuroscience*, *29*(43), 13613–13620. https://doi.org/10.1523/JNEUROSCI.2041-09.2009

Marder, E., & Bucher, D. (2001). Central pattern generators and the control of rhythmic movements. *Current Biology*, *11*(23), R986–R996. https://doi.org/10.1016/S0960-9822(01)00581-4

Marshall, T. R., O'Shea, J., Jensen, O., & Bergmann, T. O. (2015). Frontal eye fields control attentional modulation of alpha and gamma oscillations in contralateral occipitoparietal cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *35*(4), 1638–1647. https://doi.org/10.1523/JNEUROSCI.3116-14.2015

Michalareas, G., Vezoli, J., van Pelt, S., Schoffelen, J.-M., Kennedy, H., & Fries, P. (2016). Alpha-Beta and Gamma Rhythms Subserve Feedback and Feedforward Influences among Human Visual Cortical Areas. *Neuron*, *89*(2), 384–397. https://doi.org/10.1016/j.neuron.2015.12.018

Milstein, D. M., & Dorris, M. C. (2007). The Influence of Expected Value on Saccadic Preparation. *Journal of Neuroscience*, *27*(18), 4810–4818. https://doi.org/10.1523/JNEUROSCI.0577-07.2007

Milstein, D. M., & Dorris, M. C. (2011). The Relationship between Saccadic Choice and Reaction Times with Manipulations of Target Value. *Frontiers in Neuroscience*, *5*. https://doi.org/10.3389/fnins.2011.00122

Moakher, M. (2005). A Differential Geometric Approach to the Geometric Mean of Symmetric Positive-Definite Matrices. *SIAM Journal on Matrix Analysis and Applications*, *26*(3), 735–747. https://doi.org/10.1137/S0895479803436937

Molteni, E., Arrigoni, F., Bardoni, A., Galbiati, S., Villa, F., Colombo, K., & Strazzer, S. (2013). Bedside assessment of residual functional activation in minimally conscious state using NIRS and general linear models. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference*, *2013*, 3551–3554. https://doi.org/10.1109/EMBC.2013.6610309

Monti, M. M., Vanhaudenhuyse, A., Coleman, M. R., Boly, M., Pickard, J. D., Tshibanda, L., Owen, A. M., & Laureys, S. (2010). Willful modulation of brain activity in disorders of consciousness. *The New England Journal of Medicine*, *362*(7), 579–589. https://doi.org/10.1056/NEJMoa0905370

Moore, T., & Armstrong, K. M. (2003). Selective gating of visual signals by microstimulation of frontal cortex. *Nature*, *421*(6921), 370–373. https://doi.org/10.1038/nature01341

Nagel, M., Sprenger, A., Lencer, R., Kömpf, D., Siebner, H., & Heide, W. (2008). Distributed representations of the "preparatory set" in the frontal oculomotor system: A TMS study. *BMC Neuroscience*, *9*(1), 89. https://doi.org/10.1186/1471-2202-9-89

Näher, T., Bastian, L., Vorreuther, A., Fries, P., Goebel, R., & Sorger, B. (2025). Riemannian geometry boosts functional near-infrared spectroscopy-based brain-state classification accuracy. *Neurophotonics*, *12*(4), 045002. https://doi.org/10.1117/1.NPh.12.4.045002

Näher, T., Zhang, Y., Spadacenta, S., Pandinelli, M., Bastian, L., Thier, P., Hafed, Z. M., & Fries, P. (2023). *Primate Saccade Rhythmicity* (p. 2023.09.27.559710). bioRxiv. https://doi.org/10.1101/2023.09.27.559710

Naseer, N., & Hong, K.-S. (2015). fNIRS-based brain-computer interfaces: A review. *Frontiers in Human Neuroscience*, *9*. https://doi.org/10.3389/fnhum.2015.00003

Niessing, J., Ebisch, B., Schmidt, K. E., Niessing, M., Singer, W., & Galuske, R. A. W. (2005). Hemodynamic Signals Correlate Tightly with Synchronized Gamma Oscillations. *Science*, *309*(5736), 948–951. https://doi.org/10.1126/science.1110948

Noorani, I., & Carpenter, R. H. S. (2016). The LATER model of reaction time and decision. *Neuroscience & Biobehavioral Reviews*, *64*, 229–251. https://doi.org/10.1016/j.neubiorev.2016.02.018

Noori, F. M., Naseer, N., Qureshi, N. K., Nazeer, H., & Khan, R. A. (2017). Optimal feature selection from fNIRS signals using genetic algorithms for BCI. *Neuroscience Letters*, *647*, 61–66. https://doi.org/10.1016/j.neulet.2017.03.013

Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*, *10*(9), 424–430. https://doi.org/10.1016/j.tics.2006.07.005

Nyffeler, T., Wurtz, P., Lüscher, H.-R., Hess, C. W., Senn, W., Pflugshaupt, T., von Wartburg, R., Lüthi, M., & Müri, R. M. (2006). Repetitive TMS over the human oculomotor cortex: Comparison of 1-Hz and theta burst stimulation. *Neuroscience Letters*, *409*(1), 57–60. https://doi.org/10.1016/j.neulet.2006.09.011

Obrig, H., & Villringer, A. (2003). Beyond the Visible—Imaging the Human Brain with Light. *Journal of Cerebral Blood Flow & Metabolism*, *23*(1), 1–18. https://doi.org/10.1097/01.WCB.0000043472.45775.29

O'Shea, J., Muggleton, N. G., Cowey, A., & Walsh, V. (2004). Timing of target discrimination in human frontal eye fields. *Journal of Cognitive Neuroscience*, *16*(6), 1060–1067. https://doi.org/10.1162/0898929041502634

Otero-Millan, J., Macknik, S. L., Langston, R. E., & Martinez-Conde, S. (2013). An oculomotor continuum from exploration to fixation. *Proceedings of the National Academy of Sciences*, *110*(15), 6175–6180. https://doi.org/10.1073/pnas.1222715110

Otero-Millan, J., Schneider, R., Leigh, R. J., Macknik, S. L., & Martinez-Conde, S. (2013). Saccades during Attempted Fixation in Parkinsonian Disorders and Recessive Ataxia: From Microsaccades to Square-Wave Jerks. *PLOS ONE*, *8*(3), e58535. https://doi.org/10.1371/journal.pone.0058535

Otero-Millan, J., Serra, A., Leigh, R. J., Troncoso, X. G., Macknik, S. L., & Martinez-Conde, S. (2011). Distinctive features of saccadic intrusions and microsaccades in progressive supranuclear palsy. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *31*(12), 4379–4387. https://doi.org/10.1523/JNEUROSCI.2600-10.2011

Otero-Millan, J., Troncoso, X. G., Macknik, S. L., Serrano-Pedraza, I., & Martinez-Conde, S. (2008). Saccades and microsaccades during visual fixation, exploration, and search: Foundations for a common saccadic generator. *Journal of Vision*, *8*(14), 21–21. https://doi.org/10.1167/8.14.21

Owen, A. M., Coleman, M. R., Boly, M., Davis, M. H., Laureys, S., & Pickard, J. D. (2006). Detecting awareness in the vegetative state. *Science (New York, N.Y.)*, *313*(5792), 1402. https://doi.org/10.1126/science.1130197

Paillard, J., Hipp, J. F., & Engemann, D. A. (2024). *GREEN: A lightweight architecture using learnable wavelets and Riemannian geometry for biomarker exploration* (p. 2024.05.14.594142). bioRxiv. https://doi.org/10.1101/2024.05.14.594142

Paulmurugan, K., Vijayaragavan, V., Ghosh, S., Padmanabhan, P., & Gulyás, B. (2021). Brain–Computer Interfacing Using Functional Near-Infrared Spectroscopy (fNIRS). *Biosensors*, *11*(10), 389. https://doi.org/10.3390/bios11100389

Paus, T. (1996). Location and function of the human frontal eye-field: A selective review. *Neuropsychologia*, *34*(6), 475–483. https://doi.org/10.1016/0028-3932(95)00134-4

Peksa, J., & Mamchur, D. (2023). State-of-the-Art on Brain-Computer Interface Technology. *Sensors*, *23*(13), 6001. https://doi.org/10.3390/s23136001

Pennec, X., Fillard, P., & Ayache, N. (2006). A Riemannian Framework for Tensor Computing. *International Journal of Computer Vision*, *66*(1), 41–66. https://doi.org/10.1007/s11263-005-3222-z

Pinti, P., Dina, L. M., & Smith, T. J. (2024). Ecological functional near-infrared spectroscopy in mobile children: Using short separation channels to correct for systemic contamination during naturalistic neuroimaging. *Neurophotonics*, *11*(04). https://doi.org/10.1117/1.NPh.11.4.045004

Pion-Tonachini, L., Kreutz-Delgado, K., & Makeig, S. (2019). ICLabel: An automated electroencephalographic independent component classifier, dataset, and website. *NeuroImage*, *198*, 181–197. https://doi.org/10.1016/j.neuroimage.2019.05.026

Plichta, M. M., Heinzel, S., Ehlis, A.-C., Pauli, P., & Fallgatter, A. J. (2007). Model-based analysis of rapid event-related functional near-infrared spectroscopy (NIRS) data: A parametric validation study. *NeuroImage*, *35*(2), 625–634. https://doi.org/10.1016/j.neuroimage.2006.11.028

Prsa, M., Dicke, P. W., & Thier, P. (2010). The absence of eye muscle fatigue indicates that the nervous system compensates for non-motor disturbances of oculomotor function. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *30*(47), 15834–15842. https://doi.org/10.1523/JNEUROSCI.3901-10.2010

Puonti, O., Van Leemput, K., Saturnino, G. B., Siebner, H. R., Madsen, K. H., & Thielscher, A. (2020). Accurate and robust whole-head segmentation from magnetic resonance images for individualized head modeling. *NeuroImage*, *219*, 117044. https://doi.org/10.1016/j.neuroimage.2020.117044

Reingold, E. M., & Stampe, D. M. (2000). Saccadic inhibition and gaze contingent research paradigms. In *Reading as a perceptual process* (pp. 119–145). North-Holland/Elsevier Science Publishers. https://doi.org/10.1016/B978-008043642-5/50008-5

Rivet, B., Cecotti, H., Souloumiac, A., Maby, E., & Mattout, J. (2011). Theoretical analysis of xDAWN algorithm: Application to an efficient sensor selection in a P300 BCI. *EUSIPCO 2011 - 19th European Signal Processing Conference*, 1382–1386. https://hal.science/hal-00619997

Rivet, B., Souloumiac, A., Attina, V., & Gibert, G. (2009). xDAWN algorithm to enhance evoked potentials: Application to brain-computer interface. *IEEE Transactions on Bio-Medical Engineering*, *56*(8), 2035–2043. https://doi.org/10.1109/TBME.2009.2012869

Rossi, S., Antal, A., Bestmann, S., Bikson, M., Brewer, C., Brockmöller, J., Carpenter, L. L., Cincotta, M., Chen, R., Daskalakis, J. D., Di Lazzaro, V., Fox, M. D., George, M. S., Gilbert, D., Kimiskidis, V. K., Koch, G., Ilmoniemi, R. J., Lefaucheur, J. P., Leocani, L., … Hallett, M. (2021). Safety and recommendations for TMS use in healthy subjects and patient populations, with updates on training, ethical and regulatory issues: Expert Guidelines. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*, *132*(1), 269–306. https://doi.org/10.1016/j.clinph.2020.10.003

Ruff, C. C., Blankenburg, F., Bjoertomt, O., Bestmann, S., Weiskopf, N., & Driver, J. (2009). Hemispheric differences in frontal and parietal influences on human occipital cortex: Direct confirmation with concurrent TMS-fMRI. *Journal of Cognitive Neuroscience*, *21*(6), 1146–1161. https://doi.org/10.1162/jocn.2009.21097

Saager, R. B., & Berger, A. J. (2005). Direct characterization and removal of interfering absorption trends in two-layer turbid media. *Journal of the Optical Society of America. A, Optics, Image Science, and Vision*, *22*(9), 1874–1882. https://doi.org/10.1364/josaa.22.001874

Sato, T., Nambu, I., Takeda, K., Aihara, T., Yamashita, O., Isogaya, Y., Inoue, Y., Otaka, Y., Wada, Y., Kawato, M., Sato, M., & Osu, R. (2016). Reduction of global interference of scalp-hemodynamics in functional near-infrared spectroscopy using short distance probes. *NeuroImage*, *141*, 120–132. https://doi.org/10.1016/j.neuroimage.2016.06.054

Schall, J. D. (2002). The neural selection and control of saccades by the frontal eye field. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *357*(1424), 1073–1082. https://doi.org/10.1098/rstb.2002.1098

Scheeringa, R., & Fries, P. (2019). Cortical layers, rhythms and BOLD signals. *NeuroImage*, *197*, 689–698. https://doi.org/10.1016/j.neuroimage.2017.11.002

Scheeringa, R., Fries, P., Petersson, K.-M., Oostenveld, R., Grothe, I., Norris, D. G., Hagoort, P., & Bastiaansen, M. C. M. (2011). Neuronal Dynamics Underlying High- and Low-Frequency EEG Oscillations Contribute Independently to the Human BOLD Signal. *Neuron*, *69*(3), 572–583. https://doi.org/10.1016/j.neuron.2010.11.044

Schiff, N. D., Diringer, M., Diserens, K., Edlow, B. L., Gosseries, O., Hill, N. J., Hochberg, L. R., Ismail, F. Y., Meyer, I. A., Mikell, C. B., Mofakham, S., Molteni, E., Polizzotto, L., Shah, S. A., Stevens, R. D., Thengone, D., & and the Curing Coma Campaign and its Contributing Members. (2024). Brain-Computer Interfaces for Communication in Patients with Disorders of Consciousness: A Gap Analysis and Scientific Roadmap. *Neurocritical Care*, *41*(1), 129–145. https://doi.org/10.1007/s12028-023-01924-w

Schnakers, C., Vanhaudenhuyse, A., Giacino, J., Ventura, M., Boly, M., Majerus, S., Moonen, G., & Laureys, S. (2009). Diagnostic accuracy of the vegetative and minimally conscious state: Clinical consensus versus standardized neurobehavioral assessment. *BMC Neurology*, *9*(1), 35. https://doi.org/10.1186/1471-2377-9-35

Scholkmann, F., Metz, A. J., & Wolf, M. (2014). Measuring tissue hemodynamics and oxygenation by continuous-wave functional near-infrared spectroscopy—How robust are the different calculation methods against movement artifacts? *Physiological Measurement*, *35*(4), 717–734. https://doi.org/10.1088/0967-3334/35/4/717

Schroeder, C. E., & Lakatos, P. (2009). Low-frequency neuronal oscillations as instruments of sensory selection. *Trends in Neurosciences*, *32*(1), 9–18. https://doi.org/10.1016/j.tins.2008.09.012

Schroeter, M. L., Bücheler, M. M., Müller, K., Uludağ, K., Obrig, H., Lohmann, G., Tittgemeyer, M., Villringer, A., & von Cramon, D. Y. (2004). Towards a standard analysis for functional near-infrared imaging. *NeuroImage*, *21*(1), 283–290. https://doi.org/10.1016/j.neuroimage.2003.09.054

Scudder, C. A., Kaneko, C. S., & Fuchs, A. F. (2002). The brainstem burst generator for saccadic eye movements: A modern synthesis. *Experimental Brain Research*, *142*(4), 439–462. https://doi.org/10.1007/s00221-001-0912-9

Segraves, M. A., & Goldberg, M. E. (1987). Functional properties of corticotectal neurons in the monkey's frontal eye field. *Journal of Neurophysiology*, *58*(6), 1387–1419. https://doi.org/10.1152/jn.1987.58.6.1387

Shu, Z., Wu, J., Li, H., Liu, J., Lu, J., Lin, J., Liang, S., Wu, J., Han, J., & Yu, N. (2023). fNIRS-based functional connectivity signifies recovery in patients with disorders of consciousness after DBS treatment. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*, *147*, 60–68. https://doi.org/10.1016/j.clinph.2022.12.011

Si, J., Yang, Y., Xu, L., Xu, T., Liu, H., Zhang, Y., Jing, R., Li, J., Wang, D., Wu, S., & He, J. (2023). Evaluation of residual cognition in patients with disorders of consciousness based on functional near-infrared spectroscopy. *Neurophotonics*, *10*(2), 025003. https://doi.org/10.1117/1.NPh.10.2.025003

Silvanto, J., Lavie, N., & Walsh, V. (2006). Stimulation of the human frontal eye fields modulates sensitivity of extrastriate visual cortex. *Journal of Neurophysiology*, *96*(2), 941–945. https://doi.org/10.1152/jn.00015.2006

Soekadar, S. R., Kohl, S. H., Mihara, M., & Von Lühmann, A. (2021). Optical brain imaging and its application to neurofeedback. *NeuroImage: Clinical*, *30*, 102577. https://doi.org/10.1016/j.nicl.2021.102577

Sparks, D. L. (2002). The brainstem control of saccadic eye movements. *Nature Reviews. Neuroscience*, *3*(12), 952–964. https://doi.org/10.1038/nrn986

Spyropoulos, G., Bosman, C. A., & Fries, P. (2018). A theta rhythm in macaque visual cortex and its attentional modulation. *Proceedings of the National Academy of Sciences*, *115*(24). https://doi.org/10.1073/pnas.1719433115

Stanton, G. B., Goldberg, M. E., & Bruce, C. J. (1988). Frontal eye field efferents in the macaque monkey: II. Topography of terminal fields in midbrain and pons. *Journal of Comparative Neurology*, *271*(4), 493–506. https://doi.org/10.1002/cne.902710403

Strangman, G. E., Li, Z., & Zhang, Q. (2013). Depth Sensitivity and Source-Detector Separations for Near Infrared Spectroscopy Based on the Colin27 Brain Template. *PLOS ONE*, *8*(8), e66319. https://doi.org/10.1371/journal.pone.0066319

Tachtsidis, I., & Scholkmann, F. (2016). False positives and false negatives in functional near-infrared spectroscopy: Issues, challenges, and the way forward. *Neurophotonics*, *3*(3), 031405. https://doi.org/10.1117/1.NPh.3.3.031405

Tak, S., & Ye, J. C. (2014). Statistical analysis of fNIRS data: A comprehensive review. *NeuroImage, Celebrating 20 Years of Functional Near Infrared Spectroscopy (fNIRS)*, *85*, 72–91. https://doi.org/10.1016/j.neuroimage.2013.06.016

Takatoh, J., Prevosto, V., Thompson, P. M., Lu, J., Chung, L., Harrahill, A., Li, S., Zhao, S., He, Z., Golomb, D., Kleinfeld, D., & Wang, F. (2022). The whisking oscillator circuit. *Nature*, *609*(7927), 560–568. https://doi.org/10.1038/s41586-022-05144-8

Tam, N. D., & Zouridakis, G. (2015). Differential temporal activation of oxy- and deoxy-hemodynamic signals in optical imaging using functional near-infrared spectroscopy (fNIRS). *BMC Neuroscience*, *16*(Suppl 1), P245. https://doi.org/10.1186/1471-2202-16-S1-P245

Tatler, B. W., Brockmole, J. R., & Carpenter, R. H. S. (2017). LATEST: A model of saccadic decisions in space and time. *Psychological Review*, *124*(3), 267–300. https://doi.org/10.1037/rev0000054

Taylor, P. C. J., Nobre, A. C., & Rushworth, M. F. S. (2007). FEF TMS Affects Visual Cortical Activity. *Cerebral Cortex*, *17*(2), 391–399. https://doi.org/10.1093/cercor/bhj156

Teasdale, G., & Jennett, B. (1974). Assessment of coma and impaired consciousness. A practical scale. *Lancet (London, England)*, *2*(7872), 81–84. https://doi.org/10.1016/s0140-6736(74)91639-0

Thibaut, A., Schiff, N., Giacino, J., Laureys, S., & Gosseries, O. (2019). Therapeutic interventions in patients with prolonged disorders of consciousness. *The Lancet. Neurology*, *18*(6), 600–614. https://doi.org/10.1016/S1474-4422(19)30031-6

Thickbroom, G. W., Stell, R., & Mastaglia, F. L. (1996). Transcranial magnetic stimulation of the human frontal eye field. *Journal of the Neurological Sciences*, *144*(1), 114–118. https://doi.org/10.1016/S0022-510X(96)00194-3

Thielscher, A., Antunes, A., & Saturnino, G. B. (2015). Field modeling for transcranial magnetic stimulation: A useful tool to understand the physiological effects of TMS? *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 222–225. https://doi.org/10.1109/EMBC.2015.7318340

Thompson, K. G., Biscoe, K. L., & Sato, T. R. (2005). Neuronal Basis of Covert Spatial Attention in the Frontal Eye Field. *Journal of Neuroscience*, *25*(41), 9479–9487. https://doi.org/10.1523/JNEUROSCI.0741-05.2005

Tian, X., Yoshida, M., & Hafed, Z. M. (2016). A Microsaccadic Account of Attentional Capture and Inhibition of Return in Posner Cueing. *Frontiers in Systems Neuroscience*, *10*. https://doi.org/10.3389/fnsys.2016.00023

Tuzel, O., Porikli, F., & Meer, P. (2006). Region Covariance: A Fast Descriptor for Detection and Classification. In A. Leonardis, H. Bischof, & A. Pinz (Eds.), *Computer Vision – ECCV 2006* (pp. 589–600). Springer. https://doi.org/10.1007/11744047_45

Verdière, K. J., Roy, R. N., & Dehais, F. (2018). Detecting Pilot's Engagement Using fNIRS Connectivity Features in an Automated vs. Manual Landing Scenario. *Frontiers in Human Neuroscience*, *12*. https://doi.org/10.3389/fnhum.2018.00006

Villringer, A., Planck, J., Hock, C., Schleinkofer, L., & Dirnagl, U. (1993). Near infrared spectroscopy (NIRS): A new tool to study hemodynamic changes during activation of brain function in human adults. *Neuroscience Letters*, *154*(1), 101–104. https://doi.org/10.1016/0304-3940(93)90181-J

Virtanen, J., Noponen, T. E. J., & Meriläinen, P. (2009). Comparison of principal and independent component analysis in removing extracerebral interference from near-infrared spectroscopy signals. *Journal of Biomedical Optics*, *14*(5), 054032. https://doi.org/10.1117/1.3253323

Vorreuther, A., Bastian, L., Benitez Andonegui, A., Evenblij, D., Riecke, L., Lührs, M., & Sorger, B. (2023). It takes two (seconds): Decreasing encoding time for two-choice functional near-infrared spectroscopy brain-computer interface communication. *Neurophotonics*, *10*(4). https://doi.org/10.1117/1.NPh.10.4.045005

Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation Through Neural Population Dynamics. *Annual Review of Neuroscience*, *43*, 249–275. https://doi.org/10.1146/annurev-neuro-092619-094115

Wang, L., Zhang, J., Zhou, L., Tang, C., & Li, W. (2015). Beyond Covariance: Feature Representation with Nonlinear Kernel Matrices. *2015 IEEE International Conference on Computer Vision (ICCV)*, 4570–4578. https://doi.org/10.1109/ICCV.2015.519

Wassenaar, E. B., & Van den Brand, J. G. H. (2005). Reliability of Near-Infrared Spectroscopy in People With Dark Skin Pigmentation. *Journal of Clinical Monitoring and Computing*, *19*(3), 195–199. https://doi.org/10.1007/s10877-005-1655-0

Wischnewski, M., & Schutter, D. J. L. G. (2015). Efficacy and Time Course of Theta Burst Stimulation in Healthy Humans. *Brain Stimulation*, *8*(4), 685–692. https://doi.org/10.1016/j.brs.2015.03.004

Ye, J. C., Tak, S., Jang, K. E., Jung, J., & Jang, J. (2009). NIRS-SPM: Statistical parametric mapping for near-infrared spectroscopy. *NeuroImage*, *44*(2), 428–447. https://doi.org/10.1016/j.neuroimage.2008.08.036

Yeo, B. T. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Smoller, J. W., Zöllei, L., Polimeni, J. R., Fischl, B., Liu, H., & Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, *106*(3), 1125–1165. https://doi.org/10.1152/jn.00338.2011

Yger, F., Berar, M., & Lotte, F. (2017). Riemannian Approaches in Brain-Computer Interfaces: A Review. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *25*(10), 1753–1762. IEEE Transactions on Neural Systems and Rehabilitation Engineering. https://doi.org/10.1109/TNSRE.2016.2627016

Yger, F., Lotte, F., & Sugiyama, M. (2015). Averaging covariance matrices for EEG signal classification based on the CSP: An empirical study. *2015 23rd European Signal Processing Conference (EUSIPCO)*, 2721–2725. https://doi.org/10.1109/EUSIPCO.2015.7362879

Yücel, M. A., Anderson, J. E., Rogers, D., Hajirahimi, P., Farzam, P., Gao, Y., Kaplan, R. I., Braun, E. J., Mukadam, N., Duwadi, S., Carlton, L., Beeler, D., Butler, L. K., Carpenter, E., Girnis, J., Wilson, J., Tripathi, V., Zhang, Y., Sorger, B., … Boas, D. A. (2024). Quantifying the Impact of Hair and Skin Characteristics on Signal Quality with Practical Recommendations for Improvement. *bioRxiv*, 2024.10.28.620644. https://doi.org/10.1101/2024.10.28.620644

Yücel, M. A., Luke, R., Mesquita, R. C., Von Lühmann, A., Mehler, D. M. A., Lührs, M., Gemignani, J., Abdalmalak, A., Albrecht, F., De Almeida Ivo, I., Artemenko, C., Ashton, K., Augustynowicz, P., Bajracharya, A., Bannier, E., Barth, B., Bayet, L., Behrendt, J., Khani, H. B., … Zemanek, V. (2025). fNIRS reproducibility varies with data quality, analysis pipelines, and researcher experience. *Communications Biology*, *8*(1), 1149. https://doi.org/10.1038/s42003-025-08412-1

Zhai, H., Zhang, H., Li, P., & Zhang, L. (2021). Hyperspectral Image Clustering: Current achievements and future lines. *IEEE Geoscience and Remote Sensing Magazine*, *9*(4), 35–67. IEEE Geoscience and Remote Sensing Magazine. https://doi.org/10.1109/MGRS.2020.3032575

Zhang, Q., Strangman, G. E., & Ganis, G. (2009). Adaptive filtering to reduce global interference in non-invasive NIRS measures of brain activation: How well and when does it work? *NeuroImage*, *45*(3), 788–794. https://doi.org/10.1016/j.neuroimage.2008.12.048

Zhang, X., Noah, J. A., & Hirsch, J. (2016). Separation of the global and local components in functional near-infrared spectroscopy signals using principal component spatial filtering. *Neurophotonics*, *3*(1), 015004. https://doi.org/10.1117/1.NPh.3.1.015004

Zhang, Y., Brooks, D. H., Franceschini, M. A., & Boas, D. A. (2005). Eigenvector-based spatial filtering for reduction of physiological interference in diffuse optical imaging. *Journal of Biomedical Optics*, *10*(1), 11014. https://doi.org/10.1117/1.1852552

Appendix A: Data Management

Ethical approval

The research presented in this thesis involved both human participants and non-human primates across multiple institutions and was conducted in full compliance with all applicable ethical regulations.

The fNIRS experiments described in Chapter 2 and Chapter 3 were conducted at Maastricht University. Both studies were approved by the Ethics Review Committee Psychology and Neuroscience (ERCPN) of the Faculty of Psychology and Neuroscience at Maastricht University. All participants provided written informed consent and received financial compensation or participation credit hours for their participation.

The study described in Chapter 4 involved human participants and two primate species across three sites. The human experiment was approved by the ethics committee of the Faculty of Medicine at Goethe University Frankfurt (Ethikkommission des Fachbereichs Medizin der Goethe-Universität, Nr 2022-10311-). Participants gave written informed consent before participating. All procedures involving macaque monkeys were carried out in compliance with the German and European law for the protection of animals (EU Directive 2010/63/EU for animal experiments) and were approved by the Regierungspräsidium Darmstadt (permit numbers F149-1008 and F149-1007). Housing conditions complied with the same regulations, and animals were offered dietary variety, habitat enhancements, and interactive objects to support their wellbeing. All procedures involving marmoset monkeys were approved by the Regierungspräsidium Tübingen (Abteilung Tierschutz) and likewise fully complied with German and European animal protection law (EU Directive 2010/63/EU).

The study described in Chapter 5 was conducted across two sites (Mainz and Frankfurt am Main). The study was approved by both the ethics committee of the Landesärztekammer Rheinland-Pfalz, and the ethics committee of the University Hospital Frankfurt am Main. All participants were screened for MRI and TMS contraindications by an independent physician prior to inclusion.

Findability and accessibility

The fNIRS data presented in Chapters 2 and 3 were collected at Maastricht University, The Netherlands. The data and relevant metadata have been archived at Maastricht University. The analysis code associated with Chapter 2 has been made publicly available as part of the open-source Python package *pyriemann* (https://pyriemann.readthedocs.io/en/latest/).

The human eye-tracking data presented in Chapter 4 were collected at the Ernst Strüngmann Institute (ESI) for Neuroscience in Cooperation with the Max Planck Society in Frankfurt am

Main, Germany. The macaque data were collected at ESI, and the marmoset data were collected at the Werner Reichardt Centrum für Integrative Neurowissenschaften (CIN) in Tübingen, Germany. All data and relevant metadata have been archived at the respective institutes. The human eye-tracking and MRI data presented in Chapter 5 were collected at the Neuroimaging Center, Johannes Gutenberg University Mainz, and at ESI Frankfurt. These data have been archived at the respective institutions.

Interoperability and reusability

For all studies, long-lived file formats have been used to ensure that data remain usable in the future.

Data analysis scripts were written in Python and MATLAB. For Chapter 2, the Riemannian classification pipeline uses *pyriemann* , *MNE-Python*, and *scikit-learn*.

Privacy

The privacy of all human participants in this thesis has been protected using pseudonymized subject codes. A pseudonymization key linking these codes to personal data was stored separately from the research data on restricted-access storage at the respective institutions, accessible only to project members who required access in their role within the project.

For Chapters 2 and 3, the fNIRS data do not contain directly identifiable information.

For the human experiments in Chapters 4 and 5, the eye-tracking and behavioral data do not contain directly identifiable information.

A

Appendix B: Curriculum Vitae

Tim Näher was born on January 17th 1993 in Neunkirchen, Germany. After completing secondary education at the Max Planck Gymnasium Trier, he began his undergraduate studies at Troy University, United States, where he obtained a Bachelor of Science in Psychology with a Minor in Applied Behavior Analysis. He obtained a second Bachelor of Science in Psychology from Trier University, Germany.

He then moved to Maastricht University, The Netherlands, where he completed the Research Master in Clinical and Cognitive Neuroscience in 2021, graduating *summa cum laude*. His master's thesis, supervised by Prof. Peter de Weerd, received the WSUM Master Thesis Prize for the best thesis at the Faculty of Psychology and Neuroscience.

In October 2021, Tim began his doctoral research under the supervision of Prof Pascal Fries, enrolled at the Donders Graduate School of the Radboud University Nijmegen. His research was carried out at the Ernst Strüngmann Institute for Neuroscience in Cooperation with the Max Planck Society in Frankfurt am Main (2021–2024) and subsequently at the Max Planck Institute for Biological Cybernetics in Tübingen (2024–2026). During his PhD, he developed Riemannian geometry-based methods for fNIRS brain-state classification, investigated the temporal dynamics of saccadic eye movements across primate species, and contributed to various other projects in the fields of sleep, memory, and modeling. Alongside his doctoral work, he taught statistics in the interdisciplinary Master's program in Neuroscience at Goethe University Frankfurt.

From May to November 2025, Tim completed a research internship at Meta Reality Labs, developing novel deep learning architectures for real-time neuromotor decoding from EMG signals. This work resulted in a patent application for novel human-computer interaction methods. In June 2026, he joined Meta Reality Labs as a Research Scientist, working on peripheral neural interfaces for next-generation wearable devices.

Appendix C: List of Publications

Publications included in this thesis

1. **Näher, T.**, Bastian, L., Vorreuther, A., Fries, P., Goebel, R., & Sorger, B. (2025). Riemannian geometry boosts functional near-infrared spectroscopy-based brain-state classification accuracy. *Neurophotonics, 12*(4), 045002.
2. Bastian, L*., **Näher, T*.**, Vorreuther, A., Lührs, M., Benitez Andonegui, A., Fries, P., Riecke, L., & Sorger, B. (2025). Sensitive and specific fNIRS-based approach for awareness detection in disorders of consciousness: proof of principle in healthy adults. *Neurophotonics, 12*(4), 045001.*contributed equally
3. **Näher, T***., Zhang, Y*., Spadacenta, S., Pandinelli, M., Bastian, L., Thier, P., Hafed, Z. M., & Fries, P. (2023). Primate saccade rhythmicity [Preprint]. doi:10.1101/2023.09.27.559710v4 *contributed equally

Other Publications

4. Vatheuer, C. C., Dzionsko, I., Maier, S., Näher, T., van Zutphen, L., Sprenger, A., Jacob, G. A., Arntz, A., & Domes, G. (2021). Looking at the bigger picture: Cortical volume, thickness and surface area characteristics in borderline personality disorder with and without posttraumatic stress disorder. Psychiatry Research – Neuroimaging, 311, Article 111283.
5. Bastian, L., Kurz, E.-M., Näher, T., Zinke, K., Friedrich, M., & Born, J. (2024). Long-term memory formation for voices during sleep in three-month-old infants. Neurobiology of Learning and Memory, 215, Article 107987. https://doi.org/10.1016/j.nlm.2024.107987

Full List of Talks and Posters

Google scholar link

Appendix D: AI Disclosure

During the preparation of this thesis, the author used large language models for text editing, translation, paraphrasing, shortening, or summarizing. No personal, identifying, or confidential data were entered into these models. These tools were not used for generating scientific ideas, analyses, results, interpretations, or original argumentation. All AI-assisted output was critically reviewed and verified by the author, who takes full responsibility for the final content of this thesis.

Donders Graduate School

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 Behavior 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 de Donders Thesis Series. This series currently includes over 750 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 defenses please visit: http://www.ru.nl/donders/graduate-school/phd/



Radboud University



Radboudumc