Unraveling Hidden Patterns of Brain Activity:

A Journey Through Hemodynamic Deconvolution in Functional MRI

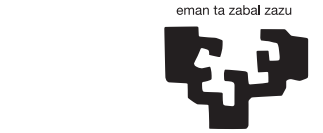

Universidad del País Vasco
Euskal Herriko Unibertsitatea

Doctoral dissertation by:
Eneko Uruñuela

Supervised by:
Dr. César Caballero-Gaudes and Dr. Miguel Ángel Veganzones

January 2024

This work was supported by the Spanish Ministry of Economy and Competitiveness (RYC-2017-21845), the Basque Government (BERC 2018-2021, PIB_2019_104, PRE_2020_2_0227), and the Spanish Ministry of Science, Innovation and Universities (PID2019-105520GB-100), and was developed at the Basque Center on Cognition, Brain and Language.

The journey is the reward.

Abstract

Functional magnetic resonance imaging data analysis is often directed to identify and disentangle the neural processes that occur in different brain regions during task or at rest, and employs the blood oxygenation-level dependent (BOLD) signal of fMRI as a proxy for neuronal activity mediated through neurovascular coupling. The goal of this thesis is to enhance and expand techniques for identifying and analyzing individual trial event-related BOLD responses based on the Paradigm Free Mapping (PFM) algorithm, which utilizes a linear hemodynamic response model and relies on regularized least squares estimators to deconvolve the neuronal-related signal that drives the BOLD effect. Notably, these techniques estimate neuronal-related activity without relying on prior paradigm information.

First of all, this work aims to establish the equivalence between the synthesis-based PFM and analysis-based Total Activation algorithms. Then, the thesis addresses the challenge of selecting the regularization parameter. This was accomplished by employing the stability selection procedure, which provides a measure of the likelihood that the estimated neuronal-related events are accurate. Building upon this, the next goal of this work is to extend the original univariate PFM formulation to a multivariate context, enabling the incorporation of spatial information through regularization terms such as the $\ell_1 + \ell_{2,1}$ mixed-norm regularization. Expanding further, the third objective of this thesis is to extend the multivariate PFM formulation to a multi-subject framework, facilitating the estimation of shared and individualized neuronal-related activity across subjects. This approach proved particularly suitable for naturalistic fMRI experiments. Lastly, the fourth and last goal of this work is to introduce an additional regularization term, the nuclear norm, into the multivariate PFM formulation. This term was employed to estimate global fluctuations during the deconvolution process and mitigate their impact on the estimation of neuronal-related activity, thereby reducing bias.

The techniques presented in this thesis were thoroughly evaluated using

simulations and experimental fMRI datasets. Comparisons were made with established fMRI analysis methods, including the single-trial general linear model, previous PFM algorithms, and other state-of-the-art techniques. Notably, the developed methods demonstrated the ability to accurately detect single trial BOLD responses in resting-state and naturalistic fMRI data, without relying on prior event information. Additionally, the potential application of multi-subject PFM in identifying both shared and individualized neuronal-related activity in more ecological datasets was explored, yielding promising results. Moreover, the utilization of low-rank and sparse PFM facilitated the extraction of global fluctuations, such as the global signal, physiological fluctuations, and motion artifacts, thereby reducing their influence on the analysis. In summary, this work demonstrates that PFM techniques can be used to reliably retrieve the neuronal-related activity from fMRI data without any prior information about the experimental paradigm, and that there now exists a formulation of PFM that is suitable for potentially any experimental setting and research question.

This thesis is available as an interactive website at phd.enekourunuela.com.

Laburpena

Normalean, erresonantzia magnetiko funtzionalarekin (fMRI) neurtutako garunaren aktibitatearen estimazioa modelo lineal orokorren (GLM) bidez egiten da. Kasu honetan ikerlariek estimuluaren denbora-egituraren informazioa badutenez, estimuluaren ondoriozko erantzun hemodinamikoa modelatu eta datuei egokitu dezakete. GLM-a erro karratu txikieko erregresio bidez ebatz daiteke zarata independiente eta identikoki banatutako (i.i.d.) Gaussiarra dela suposatuz. Hala, esperimentuarekin lotutako aktibitate neuronalaren estimazioa modelatutako erantzunaren eta neurtutako seinalearen arteko diferentzia erro karratuen batura minimizatuz lortzen da. Maiz, erregresoreen kopurua datu kopurua baino askoz txikiagoa da, eta, beraz, GLM-a problema ebatz daiteke erregularizaziorik eta suposiziorik egin gabe. Kontuan izan behar da, GLM-ak estimuluaren denbora-puntuetan besterik ez duela aktibitate neuronalaren estimaziorik egiten, eta, beraz, ez duela BOLD erantzun osoaren dinamikarik kalkulatzen.

Hala ere, GLM-a ezin da kasu guztietan erabili. Adibidez, atseden egoerako esperimentuetan erabiltzea ezinezkoa da, ez baitago estimulurik, eta, beraz, erantzun hemodinamikoaren modeloa egiteko denbora egiturarik ere ez. Teknika honen interesa nabarmen handitu da azken hamarkadan, eta horrek datuetan oinarritzen diren metodo berrien garapena ekarri du atseden egoeran, hau da, parte-hartzaileak eskanerrean etzanda eta ezer egiterik gabe egon behar duen egoeran.

Adibidez, konektibitate funtzionala (FC) garunaren area ezberdinen arteko korrelazioa neurtzen duen teknika da, eta eskala handiko garun-sareak baita azpiegiturak ikertzeko erabiltzen da. FCak pertsonak identifikatzeko ere balio du, non pertsona bakoitzaren ezaugarriei "hatz-aztarna" deitzen zaien. Hatz-aztarna hauek kognizioarekin, pertsonalitatearekin, adinarekin eta gaixotasun fenotipoekin lotuta daudela ikusi da. Konektibitate funtzional dinamikoak (dFC) konektibitatea denborarekin nola aldatzen den aztertzen du, mugitzen

den leiho bat eta taldekatze teknikak erabiliz. Fase sinkroniak (PS) ere denbora puntu bakoitzean garunaren area ezberdinen arteko sinkronia neurtzen du. Koaktibazio patroiak (CAP) garunaren alde ezberdinetan puntu bakoitzean egon daitezken seinalearen puntu maximo edo minimoak aztertzen ditu. CAP-an oinarritutako analisia innobazio koaktibazio patroien (iCAP) analisira ere luzatu da, dekonboluzioaren bidez garunaren aktibitatea eta funtzio ezberdinen azpiegiturak aztertzeko. Antzeko ideak jarraituz, "edge-centric" teknikak konektibitate funtzionalaren matrizea denbora-puntutan deskonposatzen du "edge functional connectivity" konektibitate funtzionala (eFC) neurtzeko. Hau da, eFC-ak garunaren area ezberdinek arteko ko-fluktuazioak neurtzen ditu. Patroi ia-periodikoek (QPP) ere garunaren aktibitatearen fluktuazio globalak neurtzen dituzte eta garun-sareetan aktibazio eta desaktibazio egoerak aztertzeko erabili daitezke. Azkenik, dekonboluzio teknikak gai dira erantzun hemodinamikoak lausotutako seinale garbitu eta gertaera neuronal aztertzeko erabili daitezke, gertaera hauek noiz eta non gertatzen den jakin gabe.

Tesi honen helburua fMRI datuetan hemodinamikaren dekonboluzioa egiteko metodo berriak garatu eta aplikatzea da, atseden egoeran, paradigma naturalistikoetan eta kasu klinikoetan erabiltzeko duten balioa azpimarratuz. Hortarako, hemodinamikaren dekonboluzioaren teoria aurkezten du tesiaren lehen atalak, analisi eta sintesi formulazioetan oinarritutako metodoak aztertuz. Atal honetan Paradigm Free Mapping eta Total Activation metodoak aurkeztu eta praktikan berdinak diren aztertzen da. Gero, hemodinamikaren dekonboluzio metodoen funtzionamendu egokiarentzako hain garraintzitsua den erregularizazio parametroaren aukeraketa errezten duen estabilitate-selekzio metodoa aurkeztu eta aplikatzen du hurrengo atalak. Ondoren, Paradigm Free Mapping metodoaren aldagai bakarreko formulazioa garun osora luzaten du 4. kapituluak aldagai multipleko formulazio berria aurkeztuz. Estabilitate-selekzio metodoa ere aplikatzen du atal honek. 5. kapituluak hainbat pertsonaren garunak batera aztertzeko Paradigm Free Mappingen aldagai multipleko formulazioa erabiltzen du. Formulazio hau film bera ikusi duten pertsonen garunen aktibitatearen estimazioa egiteko erabiltzen da eta estimazioaren emaitzak filmaren edukiekin lotzen ditu atal honek. Azkenik, 6. kapituluak hemodinamikaren dekonboluzio metodo berri bat aurkezten du, low-rank eta sparse deskomposizio algoritmo baten bidez. Metodo hau mugimenduak edo prozesu fisiologikoak eragindako seinale globala estimatzeko gai da, eta, ondorioz, jatorri neuronaleko gertaera hemodinamikoen estimazio garbiagoa egiten du.

2. Kapitulua: Synthesis-Based and Analysis-Based Deconvolution

Erantzun hemodinamikoaren dekonboluzioa pauso garrantzitsua da fMRI-rekin neurtutako garunaren aktibitatearen denbora eskala laburrekin lan egiteko. Nahiz eta denkonboluzio algoritmo konbentzionalak aspalditik egon diren (adibidez Wiener dekonboluzioa), eskasia bilatzen duten erregularizazioa erabiltzen duten metodo berriak fMRIarekin garunaren dinamika eta konexioak aztertzeko

interes handia sortzen ari dira. Kapitulu honetan bi metodoren oinarrizko kontzeptuak aztertzen dira, Paradigm Free Mapping eta Total Activation, modu ulergarrienan. Formulazioan ageri diren ezberdintasun guztien arren, bi metodoak teorian berdinak dira, izan ere, problema beraren sintesi eta analisi aldeak dira. Kapituluak bi metodoen berdintasun praktikoa erakusten du, simulazioen (seinale-zarata ratio ezberdinekin) eta atseden eta motore fMRI esperimentuen bidez. Emaitza berdinak ematen dituzten parametroak ebaluatzen ditu, eta metodo hauen potentziala beste metodo konbentzional batzuekin alderatzen du. Kapitulu hau metodo hauen erabiltzaileentzat erabilgarria da, bereziki dekonboluzio hemodinamikoa egiteko metodo berriak ulertzeko interesa dutenentzat, eta metodo hauen arteko ezberdintasunei buruzko galderak erantzun nahi ditu.

3. Kapitulua: Stability-Based Sparse Paradigm Free Mapping

Aktibitate neuronala fMRI datuetatik kalkulatu daiteke odolaren oxigenazio mailaren (BOLD) aldaketak noiz gertatzen diren jakin gabe, karratu txikieneko erregresio erregularizatua erabiltzen duten dekonboluzio metodoen bidez. Lan honek Sparse Paradigm Free Mapping (SPFM) dekonboluzio algoritmoarentzat bi hobekuntza proposatzen ditu: aktibitate neuronal luzea duten gertaerak detektatzeko formulazio berria; eta parametro erregularizatzaileak aukeratzeko beharra kentzen duen estabilitate-selekzio metodoa. Proposatutako metodoa fMRI datuetan ebaluatu da, eta SPFM algoritmo originalarekin eta modelo lineal orokorrarekin (GLM) konparatu da, azken hau aktibitate neuronalaren denbora modeloaren ezaguna izanda. Erakusten da estabilitate-selekzio metodoa erabiliz lortutako emaitzek GLM analisiarekin antza gehiago dutela eta SPFM algoritmo originalarekin konparatuta, aktibitate neuronalaren detekzioa nabarmenki hobetzen duela, bereziki kontraste-zarata ratio txikiko egoeretan.

4. Kapitulua: Whole-Brain Multivariate Deconvolution for Multi-Echo Functional MRI

Normalean, fMRI datuen analisiak paradigma esperimentalaren informazioa erabiltzen du garunaren aktibitatearen hipotesi-mudeloak ezartzeko. Hala ere, informazio hau okerra, osatu gabea edo eskuraezina izan daiteke hainbat kasutan, adibidez, atseden egoeretan, paradigma naturalistikoetan, edota baldintza klinikoetan. Kasu hauetan, aktibitate neuronalaren estimazioa paradigma gabeko analisi metodoen bidez lortu daiteke, adibidez hemodinamikaren dekonboluzioaren bidez. Halere, egungo hemodinamikaren dekonboluzioaren formulazioek hiru arazo edo muga garrantzitsu dituzte: 1) beraien eraginkortasuna erregularizazio parametroaren aukeraketa egokian oinarritzen da, 2) aldagai bakarrekoa izanik, garun osoan dagoen informazioa ez du kontuan hartzen, eta 3) ez dute gertaera neuronal bakoitza detektatzearen ziurtasun estatistikorik ematen. Hemen, arazo edo muga hauei konponbide bat proposatzen zaie. Zehazki, kapitulu honek MvME-SPFM (multivariate multi-echo sparse

Paradigm Free Mapping) algoritmoa proposatzen du, garun osoa kontuan hartzen duen eta norma misto baten bidez informazio espaziala gehitzen duen hemodinamikaren dekonboluzio metodo berri bat. Gainera, MvME-SPFM-ek erregularizazio parametroak aukeratzeko beharra kentzen duen estabilitate-selekzio metodoa erabiltzen du, eta, halaber, aktibitate neuronalaren probabilitatea estimatzen du voxel eta denbora-puntu bakoitzean, estabilitate bideen kurba azpiko azala (AUC) kalkulatuz. Formulazio hau multi-echo fMRI datuentzako egokitua dago, eta horrek baliabide gehiago eskaintzen ditu echo-time (TE) denbora-puntuekin linealki lotutako fluktuazioak aztertuz BOLD jatorriko fluktuazioak isolatzeko; eta aktibitate neuronalaren estimazioak unitate fisiologikoki interpretagarriak izateko (adibidez, ageriko errelaxazio transbersalaren aldaketekin ΔR_2^*). Kapitulu honek MvME-SPFM algoritmoak beste hemodinamikaren dekonboluzio metodoek baino emaitza hobeak lortzen dituela erakusten du, eta, halaber, GLM analisi linealaren emaitzekin duen espazio- eta denbora-koherentzia handiagoa erakusten du, baita gertaera neuronal bakoitzaren denbora-puntu bakoitzean ere. Ondorioz, proposatutako MvME-SPFM algoritmoak aktibitate neuronalaren estimazio fidagarriagoa eskaintzen du, ΔR_2^* unitateetan, BOLD gertaerak noiz gertatzen diren jakin gabe garunaren dinamikak aztertzeko. Algoritmo hau publikoki eskuragarri dago splora Python paketearen barruan hurrengo web gunean: https://github.com/ParadigmFreeMapping/splora

5. Kapitulua: Multi-Subject Paradigm Free Mapping

fMRI bidezko paradigma naturalistikoak (esaterako, filma ikustea) are eta gehiago erabiltzen ari dira kognizio neurozientzian, ekologikoagoak diren baldintzetan kognizioa, adibidez, gertaera segmentazioa eta memoriaren eraketa, aztertzeko. Baina, egungo analisi metodoak mugatuta daude eta ez dira aktibitate neuronalaren dinamika espazio-tenporal zehatzak ondo aztertzeko gai. Adibidez, GLM analisi tradizionalak eta encoding modeloak BOLD aktibitate mapak kalkulatzeko gai dira estimuluaren ezaugarriak ezagututa, baina estimuluaren zer ezaugarri diren garrantzitsu eta hauek modu zehatz batean modelatzeko suposizio handiak egin behar dituzte. Bestalde, modelorik gabeko metodoek, adibidez, pertsonen arteko korrelazioak eta gertaera segmentazio algoritmoek, ez dute estimuluaren ezaugarririk suposatzen, baina informazioa estimuluaren zenbait denbora-puntutan konbinatzea eskatzen dute, eta, ondorioz, ez dira berreskuratutako seinalea estimuluaren denbora-puntu zehatzetara lotzeko gai. Defektu hauei aurre egiteko, kapitulu honetan metodo berri bat proposatzen da: multi-subject Paradigm Free Mapping (msPFM). Proposatutako metodoak, existitzen diren beste metodoek ez bezela, talde eta banakako aktibitate neuronalaren estimazioa egiten du erresoluzio espazial eta tenporal zehatzenean. msPFM algoritmoak hainbat aldagaieko karratu txikieko erregresio erregularizatua erabiltzen du estimuluaren denbora-egituran oinarritu gabe aktibitate neuronalaren estimazioa egiteko. Kapitulu honek, algoritmo honen eraginkortasuna simulazio eta esperimentazio fMRI datuetan ebaluatu, eta, ondorioz, estimuluaren behe eta erdi mailako ezaugarriekin lotutako aktibitate neuronala detektatzeko gai dela

erakusten du. Gainera, metodoak garunaren azalean gradiente formako konfigurazio espazio-tenporala aurkitzen du, literaturan aurkeztu diren emaitzekin bat eginez. Laburki, msPFM-ek fMRI paradigma naturalistikoetan aurki daitezken garunaren aktibitatearen dinamika espazio-tenporal konplexuak aztertzeko tresna ahaltsua eskaintzen du. msPFM algoritmoaren Python kode irekia eskuragarri dago hurrengo estekan: https://github.com/ParadigmFreeMapping/msPFM.

6. Kapitulua: Sparse and Low-Rank Paradigm Free Mapping

Mugimenduak edo prozesu fisiologikoek (esaterako, arnasketa sakonak) eragindako garun osoko seinale aldaketek egungo dekonboluzio algoritmoen eraginkortasuna murrizten dute, eta, ondorioz, jatorri neuronaleko gertaera hemodinamikoak ezin dira ondo interpretatu. Kapitulu honek garunaren seinale globala eta aktibitate neuronala batera estimatzeko metodo berri bat proposatzen du, BOLD gertaerak noiz gertatzen diren jakin gabe, low-rank eta sparse deskomposizio algoritmo baten bidez. Proposatutako metodoaren eraginkortasuna simulazio eta esperimentazio fMRI datuetan ebaluatu, eta metodo berriaren emaitzak egungo metodo konbentzional baten emaitzekin alderatu ditu lan honek, azken metodo honek aktibitate neuronala noiz gertatzen den jakinik. Emaitza hauek erakusten dutenez, proposatutako low-rank eta sparse Paradigm Free Mapping algoritmoa esperimentuarekin lotutako mugimendua azaltzen duen seinale globala estimatzeko gai da, eta, aldi berean, aktibitate neuronala zehaztasun handiz estimatzen du.

Acknowledgements

The pages of this thesis are just a tiny reflection of the experiences, learnings, skills, collaborations, and friendships that I have gained over the past five years. Since the day I joined the BCBL to do my master's thesis until today, as I write these words with the thesis written, I have been very fortunate to have been surrounded by amazing people who have helped me grow as a scientist, a researcher, a professional, and a person.

First and foremost, it is my family that I have to thank for their unconditional support and love. Aita, ama, Eñaut, mila esker hona iristeko erreztasun guztiak emateagatik eta etxeko giroarekin lana ahazten laguntzeagatik.

Speaking of taking my mind off work, mila esker kuadrila. Nahiz eta ez garen hain maiz ikusi azken urte hauetan, ikusi garen bakoitzean ordu batzuz tesiaz ahazten lagundu nauzue. Ospatu beharko dugu.

I have also met incredble people and made some friends along the way I feel very grateful for. Thank you Eleonora, Maca, Edith, Trisha, Jessi, Christophoros, Asier, Giorgio, Mina, Hana, Thomas, Soan, Kaya, Maksim, Sandra. Especial thanks to Marta for all the F1 talk and the tour around Madrid; to Irene for being a great friend and for letting me crash on your couch.; and to Alberto for being the lovely person you are and for being a great friend. I know I can always count on you.

A very special thank you to two of the most important people in my life, Chiara and Dani (and Biruji!). If I have grown as a person over the past five years it is in no small part thanks to you. Since the summer of 2019 you have been my my best friends and I cannot thank you enough for all the love and support you have given me. I am so happy to have you in my life and I hope we can keep making memories together for many years to come. I love you both.

Thank you to pals in Oxford who have made me feel at home in the UK. Thank you Joana and Ben, Becca, Gabby, Sammi, Grace and Adam, Minying, Amin, Camron, Ti, Michael. You are all beautiful people and I cannot wait to

see you again.

Thank you to the Iberian Chapter of the ISMRM for making the conferences so much fun. Special thanks to my crazy amiga iberica Teresa, and to the rest of the people who have been heavily involved in the chapter, especially to Rita, Rui, and Jose.

I would like to thank the BCBL, and especially Ana Fernández, for helping me get the grant that allowed me to pursue my PhD.

Many thanks to the members of the tedana community. I have learned a lot about multi-echo fMRI and open source software development from you. Special thanks to Elizabeth, Dan, Taylor, Josh, Logan. I owe you most of the knowledge I have about open source software development.

Thank you to everyone I have collaborated with over the years and to everyone I have hold discussions with. I have learned a lot from you and I hope you have learned something from me as well. Special thanks to Javier, Óscar, Dimitri, Patricia, Gari, Stephen, Inês and Mairi.

I would like to thank all the members of the SPiN lab, past and present, for making the BCBL a great place to work. Thank you for all the discussions, the help, the feedback, the laughs, and making work fun. Thank you Maite, Borja, Stefano, Vicente, Inés, Marco, Luca, Cristina T., Cristina C., Iñigo, Ryland, and David.

I would also like to thank Emily Finn and the rest of the Finn Lab for hosting me in the New Hampshire forest at Dartmouth College for three months. It was a dream come true to spend some time in the US, and to be able to do it with you all made it even better. I learned a lot from you and I hope you got to learn a bit from me as well. I really hope the work we did together will be something you're proud of and that it will serve the lab and the community well. Thank you to the whole lab for making me feel at home. Special thanks to Clara for being awesome at helping me with everything related to our project. Of course, I cannot forget my Korean sister, Hongji. You made my time in the US so much better and I am so happy we got to spend so much time together. I hope to see you all again soon.

I would like to thank my supervisors, Dr. Miguel Ángel Veganzones and Dr. César Caballero Gaudes.

Miguel, I feel very fortunate to have met you and to have had the opportunity to work with you. I believe I have learned quite a bit from you even in the short time we have worked together, and all I can say is I wish we had more time.

César, I'm getting emotional and I haven't even started writing. I am incredibly grateful for the chance you gave me in May-June 2018 and the faith you had in me to not only do my master's thesis with you, but also to continue with a PhD. I have to admit that you are the reason I decided to stay and do a PhD. I knew I would learn a lot from you and I was right. It is really hard for me to write this because it means that my time working with you is coming to an end and believe me that is probably going to be the hardest part of my PhD. I have learned so much from you, not only about MRI, methods, and neuroscience, but also about how to be a good leader and I hope one day I am half as good as you are. People often use the word supervisor, but you have

been more than that to me, a mentor and a friend. I hope I have been a good PhD student and that I have made you proud making all these advances to your baby Paradigm Free Mapping. Thank you so so much. You have made these five years an incredible experience and I will always be grateful to you. A big thank you to María and the kids as well, Claudia and Emma, you are all wonderful and I feel very fortunate to have met you.

Finally, I would like to thank you Gen for your love and support, and for bringing extraordinary magic to my life. This year and a half has been the most beautiful time of my life and I cannot think of a better way to finish my PhD than with you by my side, my bidelagun.

Contents

Acronyms

AIC Akaike information criterion.

ASL arterial spin labeling.

ATS activation time series.

AUC area under the curve.

BIC Bayesian information criterion.

BOLD blood oxygenation-level dependent.

CAP co-activation pattern.

CBF cerebral blood flow.

CBV cerebral blood volume.

CMRO2 cerebral metabolic rate of oxygen.

CNN convolutional neural network.

CSF cerebrospinal fluid.

DAN dorsal attention network.

dFC dynamic functional connectivity.

DMN default mode network.

ECoG electrocorticography.

EEG electroencephalography.

eFC edge-centric functional connectivity.

EROS event-related optical signaling.

FC functional connectivity.

FISTA fast iterative shrinkage-thresholding algorithm.

fMRI functional magnetic resonance imaging.

GLM general linear model.

GM gray matter.

GSBS greedy state boundary search.

HMM hidden Markov model.

HR heart rate.

HRF hemodynamic response function.

ICA independent component analysis.

iCAPs innovation-driven co-activation patterns.

ISC inter-subject correlation.

LARS least angle regression.

ME multi-echo.

ME-SPFM Multi-Echo Sparse Paradigm Free Mapping.

ME-ICA multi-echo independent component analysis.

MEG magnetoencephalography.

MFISTA-VA Monotone Fast Iterative Shrinkage-Thresholding Algorithm with Variable Acceleration.

MR magnetic resonance.

msPFM Multi-Subject Paradigm Free Mapping.

Mv-SPFM Multivariate Sparse Paradigm Free Mapping.

nFC node-centric functional connectivity.

NIRS near-infrared spectroscopy.

PET positron emission tomography.

PFM Paradigm Free Mapping.

PopSync+ population synchrony.

PPI psychophysiological interactions.

QPP quasi-periodic patterns.

ROC receiver operating characteristic.

ROI region of interest.

RSSD root sum of squares difference.

RVT respiratory volume per time.

SNR signal-to-noise ratio.

SPFM Sparse Paradigm Free Mapping.

SPLORA-PFM Sparse and Low Rank Paradigm Free Mapping.

SVD singular value decomposition.

TA Total Activation.

TCD transcranial Doppler ultrasound.

TE echo time.

TR repetition time.

WM white matter.

Chapter

1

Functional Magnetic Resonance Imaging and Blood Oxygen Level Dependent Signal

1.1 Functional Imaging of the Human Brain

For over a century, neuroscientists and neurophysiologists have dedicated their efforts to uncovering the functional organization of the brain. Through the use of various functional neuroimaging techniques with complementary temporal and spatial resolutions, they aim to reveal the neuro-anatomical localization and dynamic changes of brain activations.

While non-invasive electrophysiological techniques such as electroencephalography (EEG) or magnetoencephalography (MEG) excel in temporal resolution (10-100 ms), their spatial resolution is relatively poor (several mm or cm) (Baillet et al., 2001; Hämäläinen et al., 1993). In contrast, invasive electrophysiological techniques, including patch clamps (Neher et al., 1978), single-unit or multi-unit recordings (Arieli et al., 1995), and electrocorticography (ECoG) (Miller et al., 2007; Nir et al., 2008), offer high spatial resolution.

Functional magnetic resonance imaging (fMRI) and Positron Emission Tomography (PET) provide insight into cerebral blood flow and oxygen metabolism, indirectly associated with neural activations. With spatial resolutions in the order of millimeters, these techniques are capable of capturing activations in both cortical and deep brain structures, but their temporal resolution is limited by the sluggish dynamics of hemodynamic changes (Dale & Halgren, 2001). fMRI offers higher temporal and spatial resolution compared to PET, making it more suitable for studying the temporal responses to short neuronal-related events. However, PET has the advantage of measuring well-defined physiological quantities such as cerebral blood flow (CBF), cerebral blood volume (CBV), or cerebral metabolic rate of oxygen or CMRO2 (Fox et al., 1986; Friston et al., 1993).

Functional MR spectroscopy (MRS) is an alternative method for metabolic imaging, offering quantitative measurements of functional changes in neu-

rometabolite and neurotransmitter concentrations within a specific brain region. However, it is important to note that MRS has lower spatial and temporal resolutions compared to fMRI (Morris, 1999). Typically, MRS is performed in a single voxel with an isotropic size of a few centimeters, or in multiple voxels using chemical shift imaging.

Lastly, optical diffusion imaging techniques, such as near infrared spectroscopy (NIRS) (Kleinschmidt et al., 1996; Villringer et al., 1993), diffusion optical tomography (DOT) (White & Culver, 2010), or event-related optical signaling (EROS) (Gratton & Fabiani, 2001; Gratton et al., 1997), utilize an optical imaging device placed on the scalp to measure changes in cortical blood flow. However, due to the need for the light to penetrate through the skull, these techniques have lower spatial resolution compared to MR imaging techniques and are limited to studying the cortical surface. The temporal resolution of EROS is similar to MEG and EEG (in the order of milliseconds), while NIRS has a ten times higher temporal resolution to fMRI.A notable advantage of optical imaging is its lower cost and portability compared to other techniques such as MRI or MEG.

1.2 Blood Oxygen Level Dependent Functional MRI

In order to fully comprehend the assumptions and methods presented in this thesis, it is essential to review some basic concepts related to the physical and physiological basis of the blood oxygen-level dependent (BOLD) effect. MRI techniques offer a range of approaches to detect the increased metabolic demand associated with brain function. These include utilizing the BOLD contrast, changes in CBV using contrast agents such as gadolinium (Dean et al., 1992), ferumoxytol (Christen et al., 2012) or hyperoxic contrasts (Bulte et al., 2007), and assessing CBF through arterial spin labeling techniques (ASL) (Buxton, 2009). BOLD fMRI, introduced by Ogawa et al., 1990, 1992, is particularly advantageous as it requires no exogenous contrast agent and exhibits higher sensitivity compared to CBF-based contrast, such as ASL. Consequently, it has gained widespread usage for functional brain imaging.

1.2.1 Physiological Basis of the BOLD Constrast

The signal contrast in BOLD fMRI images arises from variations in the local magnetic susceptibility, χ, caused by disparities in blood hemoglobin oxygen concentration. As local neuronal activity intensifies, there is a rise in oxygen consumption, leading to an augmented supply of oxygenated blood that diffuses passively through the capillary blood vessels to the tissue. When oxygen binds to hemoglobin (forming oxyhemoglobin), it exhibits slight diamagnetic properties relative to the tissue, causing the molecule to repel the magnetic field. In contrast, deoxygenated hemoglobin is paramagnetic compared to the tissue, attracting the magnetic field. Consequently, the magnetic field becomes distorted near deoxygenated red blood cells, creating higher local magnetic field gradients in

the surrounding tissue, which results in spin dephasing. Reduced oxygenation amplifies the spin dephasing effect, shortening the tissue's T_2^* and diminishing the amplitude of the MR signal in T_2^*-weighted images. Conversely, higher oxygen concentration aligns the susceptibility of the blood with that of the surrounding tissue, reducing the local magnetic field gradient, increasing T_2^*, and raising the measured MR signal amplitude by a few percent. This forms the basis of BOLD fMRI, where changes in blood oxygenation serve as an intrinsic contrast mechanism in T_2^*-weighted images, enabling the identification of cortical regions exhibiting functional activity characterized by increased oxygen demand and supply (Bandettini et al., 1992; Belliveau et al., 1991; Kwong et al., 1992; Ogawa et al., 1990, 1992; Turner et al., 1991).

1.2.2 Temporal Characteristics of the BOLD Signal: the Hemodynamic Response Function

The BOLD effect does not directly reflect neuronal activity, but rather measures the hemodynamic response associated with it (Logothetis, 2008; Logothetis et al., 2001). The relationship between the hemodynamic response and the underlying neuronal activity is complex involving dynamic changes in CBF, CBV and $CMRO_2$ (Buxton, 2009). After neuronal activity increases in a brain region, there is an initial decrease in blood oxygenation due to oxygen consumption which might cause a small initial dip in the hemodynamic response following the first second after the activation (Ernst & Hennig, 1994; Menon et al., 1995). Although this initial dip is not always observed in fMRI (Behzadi & Liu, 2006; Buxton, 2009; Hu et al., 1997), it is suggested that it maps more accurately the site of neural activity (Duong et al., 2000; Malonek & Grinvald, 1996). Afterwards, the local supply of oxyhemoglobin increases more than it is strictly demanded, probably to ensure a large oxygen gradient across capillary walls so that there is a high rate of transfer of oxygen or glucose to tissue (Logothetis, 2008), generating a positive BOLD response due to an excess of oxyhemoglobin. Negative BOLD responses have also been observed associated with neuronal deactivations (Shmuel et al., 2006). Regardless of the polarity of the response, the BOLD response peaks between 5 and 8 s after the activation starts and its amplitude depends on the type of stimulus and the magnetic field strength. For instance, for visual stimulation the signal change is 2-3 % at 1.5T, 4-6 % at 3T, 7-10 % at 7T (van der Zwaag et al., 2009). After the stimulus ceases, there is a return of the BOLD response to baseline, often followed by a post-stimulus undershoot due to an increase of deoxyhemoglobin which may last for several seconds until the response returns to baseline. The cause of the post-stimulus undershoot is also not completely understood, whether this is a vascular, neural or metabolic effect (Buxton, 2009; Chen & Pike, 2009). In summary, the time scale of the BOLD response is much slower than the time scale of neural activity and the return of the BOLD signal to baseline level after a short stimulus may occur more than 30 s from the stimulus onset. The temporal characteristics of the BOLD response are usually modelled by a hemodynamic response function (HRF). Figure 1.1 shows the shape of typical HRF, along with the initial dip

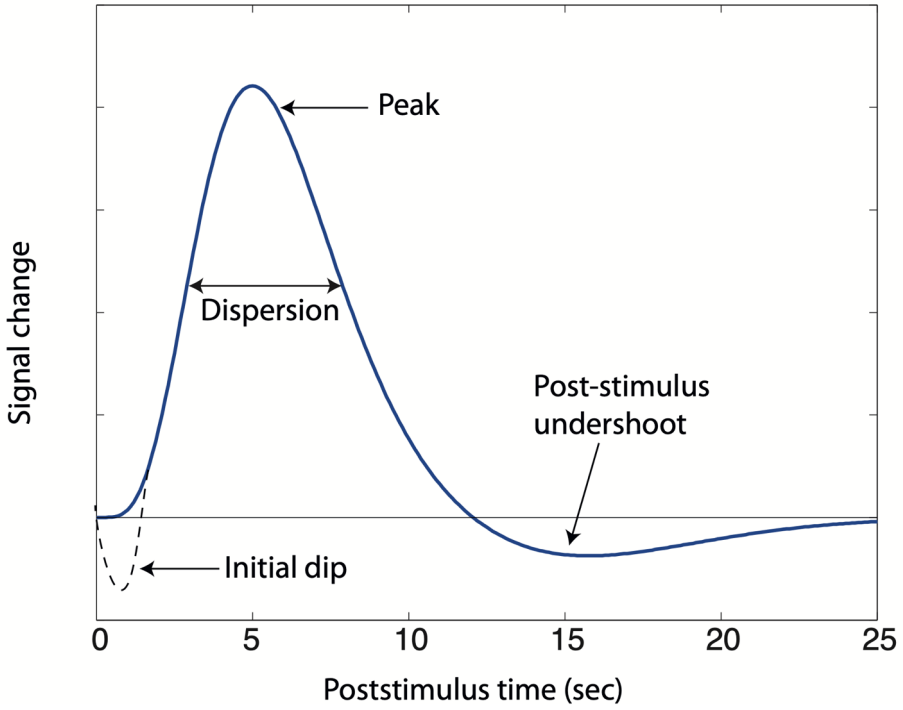

Figure 1.1: Temporal characterization of the haemodynamic response function. The shape in bold line corresponds to the canonical haemodynamic response function.

for illustration of this effect. In this figure, the HRF plotted is the well-known canonical HRF (Friston, 2007; Friston et al., 1998), which is defined as the difference of two-gamma functions:

$$h(t) = g(t; a_1, b_1) - \frac{1}{c} g(t; a_2, b_2), \tag{1.1}$$

where the Gamma function is given by

$$g(t; a, b) = \frac{b^a t^{a-1} e^{-bt}}{\Gamma(a)}. \tag{1.2}$$

The canonical HRF, as described in literature (Friston, 2007), is characterized by specific parameters. These parameters include a time-to-peak (a/b) of 6 s and dispersion (a/b^2) of 1 s for the initial overshoot. Additionally, it has a time-to-peak of 16 s and dispersion of 1 s for the undershoot. The overshoot-undershoot ratio (c) is approximately 6. However, it is important to note that the BOLD response exhibits variability across trials, brain regions, and subjects (Aguirre, Zarahn, & D'Esposito, 1998b; Duann et al., 2002; Handwerker et

al., 2004; McGonigle et al., 2000; Smith et al., 2005). To account for this variability, alternative HRF shapes have been proposed in the scientific literature. These include single gamma functions (Boynton et al., 1996; Cohen, 1997), two-gamma functions with different HRF parameters (Glover, 1999), Poisson functions (Friston, Jezzard, & Turner, 1994). and Gaussian functions (Kruggel & von Cramon, 1999; Rajapakse et al., 1998).

1.3 Noise in BOLD fMRI

In addition to neuronal-related activity, the BOLD signal presents multiple sources of noise related to hardware-related artifacts and drifts, head motion, confounding physiological fluctuations (Bianciardi et al., 2009; Jorge et al., 2013) as well as image distortions related to data acquisition that should be accounted for and corrected during fMRI data preprocessing.

1.3.1 Motion and Other Artifacts

For instance, the images could present geometric distortions in the phase direction of the acquisition that can be removed by means of mapping field distorsions, either using field maps or using two images acquired in opposite phase-encoding directions, and then with a non-linear transformation of the fMRI volumes (e.g. TOPUP, Andersson et al., 2003; Glasser et al., 2018). Also, fMRI images are typically acquired slice by slice and the difference in their time of acquisition can be compensated via slice timing correction, although this step can also introduce confounding effects in the data due to signal interpolation (Parker & Razlighi, 2019). More generally, realigning all fMRI volume to a reference image can deal with part of the artifacts introduced by head motion (Friston, Holmes, et al., 1994). However, this step does not remove the effect of motion completely (Caballero-Gaudes & Reynolds, 2017).

The most straight-forward way to deal with signal artifacts is to model them as regressors of non-interest along with the task regressors (in tIA-fMRI) or to project them out of the fMRI data if there is no task paradigm, as in resting-state (RS) experiments. For instance, motion-related effects can be expressed as a set of the relative translations and rotation obtained during realignment, considering their first derivative and their squared transformation for up to 24 regressors for a better denoising (Friston, Holmes, et al., 1994). Similarly, very low frequency trends due to hardware-related inestabilities can be modelled as a set of basis functions (e.g. using Legendre polynomials, Discrete Cosine Transform).

Furthermore, large deviations in the fMRI signal (e.g., spikes) caused by motion jerks and scanner noise can be removed through scrubbing or censoring (Power et al., 2012). This process consists in identifying those fMRI volumes characterized by abrupt changes in the BOLD signal and removing or interpolating them. The identification can be performed by using summary metrics of motion, like Framewise Displacement (Power et al., 2012), or by observing transient changes in the signal (DVARS, Power et al., 2012; Smyser

et al., 2010). However, it is important to notice that, if correctly accounted for, scrubbing could reduce the degrees of freedom in statistical analysis (Mascali et al., 2021), leading to biases in second-level analysis between subjects that move too much and others, or introduce discontinuity in the signal itself, limiting the use of particular analyses dependent on signal continuity (e.g. ICA, see Caballero-Gaudes and Reynolds, 2017), and biasing the estimation of functional connectivity (Mascali et al., 2021). Moreover, interpolating the signal could introduce spurious changes in the signal.

Another common approach to remove not only motion, but also other sources of noise, is based on data decomposition techniques. For instance, ICA can be leveraged to model, identify and remove motion artifacts as well as other sources of noise (Griffanti2014ICAbasedartifact; Behzadi et al., 2007; Muschelli et al., 2014; Pruim, Mennes, Buitelaar, & Beckmann, 2015; Pruim, Mennes, van Rooij, et al., 2015; Salimi-Khorshidi et al., 2014). Ideally, the best candidate to identify noisy timeseries would be temporal ICA, i.e., a decomposition in which the independence is forced in the temporal domain (Glasser et al., 2018; Smith et al., 2012). However, such approach is not feasible in a normal fMRI context since it would require the samples in time to be much higher than the samples in space (Smith et al., 2012). Hence, spatial ICA is the most common application for fMRI decomposition, although this might lead to detect spurious components that contain both true BOLD signal and noise (Caballero-Gaudes & Reynolds, 2017). Alternatively, several fMRI sessions could be concatenated to apply temporal ICA, although this approach would lead to the impossibility of removing session-specific noise. The challenging factor in adopting ICA for denoising is the classification of the independent components. Although manual classification is still the approach with the best outcome (Griffanti et al., 2017), it is time consuming, it requires trained researchers, and the result is dependent on the observer. For this reason, different approaches for automatic classification of ICA components have been proposed in time, from full classifiers (FIX, Salimi-Khorshidi et al., 2014) to approaches specifically targeting motion artifacts (ICA-AROMA, Pruim, Mennes, Buitelaar, and Beckmann, 2015; Pruim, Mennes, van Rooij, et al., 2015).

Alternatively, decomposing the signal of white matter (WM) and cerebrospinal fluid (CSF) into principal components and considering the first few (a technique called anatomical CompCor, see Behzadi et al., 2007) can help retrieving proxies of motion-related artifacts and physiological fluctuations (Behzadi et al., 2007; Muschelli et al., 2014). In fact, it has been shown that CompCor can be more effective in denoising motion artifacts than ICA based techniques and censoring (Mascali et al., 2021).

Noise in fMRI can also be reduced by using multi-echo (ME) acquisitions that sample the data at multiple successive echo times (TE). A weighted combination of the multiple echoes based on each voxel's T_2^* value (Posse et al., 1999) or temporal signal-to-noise ratio (Poser et al., 2006) can smear out random noise and enhance the sensitivity to the BOLD contrast. In fact, compared with single-echo data, this optimal combination (Liu et al., 2022) can improve the mapping of neuronal activity at 3T (Fernandez et al., 2017) and 7T (Puckett

et al., 2018), with results comparable to other preprocessing techniques requiring extra data such as RETROICOR (Atwi et al., 2018). Optimal combination of multiple echo volumes can also improve sensitivity, specificity, repeatability and reliability of fMRI mapping (Cohen & Wang, 2019; Cohen et al., 2021).

Furthermore, assuming a monoexponential decay, the voxelwise fMRI signal acquired at a given echo time TE, i.e. $S(TE)$, can be expressed in signal percentage change as:

$$\frac{S(TE) - \overline{S}(TE)}{\overline{S}(TE)} \approx \Delta\rho - \mathrm{TE} \cdot \Delta R_2^* + n, \tag{1.3}$$

where $\overline{S}(TE)$ is the average signal at a given TE, $\Delta\rho$ represents non-BOLD related changes in the net magnetisation, ΔR_2^* represents BOLD-related susceptibility changes (and is the inverse of ΔT_2^*), and n denotes random noise (Kundu et al., 2012, 2013). As the BOLD-related signal can be expressed as a function of the TE, whereas noise-related (i.e., non-BOLD) changes in the net magnetization are independent of TE, the information available in multiple echoes can be leveraged for the purpose of denoising. For example, in a dual-echo acquisition where the first TE is suficiently short, the first echo signal mainly captures changes in $\Delta\rho$ rather than in ΔR_2^*. It is then possible to remove artifactual effects, through voxelwise regression, from the second echo signal acquired at a longer TE with appropriate BOLD contrast (Bright & Murphy, 2013).

Collecting more echoes opens up the possibility to leverage ICA and automatically classifying independent components into BOLD-related (i.e., describing ΔR_2^* fluctuations with a linear TE-dependency) or noise (i.e., independent of TE, related to non-BOLD fluctuations in the net magnetization $\Delta\rho$), an approach known as ME independent component analysis (ME-ICA, Kundu et al., 2012, 2013, 2017). Compared to single-echo data denoising, ME-ICA can improve the mapping of task-induced activation (DuPre et al., 2021; Gonzalez-Castillo et al., 2016; Lombardo et al., 2016), for example in challenging paradigms with slow-varying stimuli (Evans et al., 2015) or language mapping and laterality (Amemiya et al., 2018). It also outperforms single-echo ICA-based denoising of resting-state fMRI data (Dipasquale et al., 2017; Lynch et al., 2020), and provide more effcient and reliable functional connectivity mapping in individual subjects (Lynch et al., 2020) and in brain regions where traditional single-echo acquisitions offer reduced signal-to-noise ratio, such as the basal forebrain (Markello et al., 2018). Finally, ME-ICA also enhances the deconvolution of neuronal-related signal changes (Caballero-Gaudes et al., 2019).

1.3.2 Physiological Artifacts

When employing BOLD fMRI as an intrinsic contrast to investigate neural correlates, it becomes essential to decouple the neurovascular coupling. In this context, physiological signals may introduce noise, requiring their modeling to account for and minimize their associated variance during preprocessing or data analyses (Caballero-Gaudes & Reynolds, 2017; Liu, 2016). The principal frequencies characterising physiological signals like cardiac pulse and respiration

are in a different band compared to those of the neural-related BOLD signal: namely, the primary component of cardiac related fluctuations are around 1 Hz, while the respiratory related ones are around 0.3 Hz. Thus, if the temporal sampling is high enough, a simple band-pass filter could easily remove their confounding effects (Biswal et al., 1995; Chuang & Chen, 2001; Lowe et al., 1998). The downside is that this approach will remove the BOLD-related signal frequencies in the same range as well, and it will not remove the impact of physiological frequencies in the same range as the BOLD signal (Caballero-Gaudes & Reynolds, 2017). This is especially true if the temporal sampling is low and the physiological signal is aliased in the BOLD-related frequency range. Moreover, physiological signal, and respiration in particular, has an impact on other sources of noise, like magnetic field perturbation (Raj et al., 2001), and motion (Fair et al., 2020; Pais-Roldán et al., 2018; Power et al., 2019) that should be taken into account and removed (Gratton et al., 2020).

An easy way to remove such perturbations is to remove the average brain signal (also called global signal), since it is often considered as a proxy of the combined impact of different sources of noise, especially related to movement or physiological in nature. However, its removal is controversial, since it can heavily alter the interpretation of BOLD fMRI (Power et al., 2017). For this reason, Power et al., 2018 proposed to decompose fMRI data in low and high rank components, and to consider the first few low rank components timeseries as noise. This technique, called GODEC, showed improved denoising of fMRI data after ME-ICA (Power et al., 2018; Zhou & Tao, 2011).

As an alternative, the average signal in the white matter (WM) and cerebrospinal fluid (CSF) can be used as a proxy of physiological noise, since no neural-related signal is present in these tissues, that are conversely dominated by cardiac pulsatility and respiration (Anderson et al., 2011; Jo et al., 2010), although more recently Attarpour et al., 2021 showed that the average CSF does not represent cardiac fluctuations properly. Specifically, ICA based decomposition can be set up to retrieve physiological-related signals, both in space (CORSICA, Perlbarg et al., 2007) and in time (PESTICA, Beall and Lowe, 2007).

An alternative to data-driven approaches consists in acquiring physiological signals such as cardiac pulse and respiration effort during the imaging session, opening up the possibility to adopt more model-based approaches to deal with physiological noise. For instance, it is possible to estimate the frequencies of the amplitude envelope of cardiac and respiratory signals, and then selectively filter them from fMRI data (Biswal et al., 1996). However, the main frequency and the harmonics of physiological fluctuations are usually aliased with the spectra of BOLD components with neurobiological relevance, and bandpass filtering would remove them as well.

Alternatively, it is possible to use the measured cardiac and respiratory signals to model their pseudo-periodic fluctuation that are phase-locked to the fMRI signal. Cardiac and respiratory phases can be estimated from signal recordings, then their Fourier expansion can be removed from the data in a slice-dependent manner at the beginning of the preprocessing, an approach known as RETROICOR, Glover et al., 2000. However, RETROICOR does not

remove completely the effect of physiological signal from the data, especially regional low frequency effects that vary between brain regions (Birn et al., 2006; Chang et al., 2009; Shmueli et al., 2007).

Noticeably, slow variations in the heart rate (HR, Shmueli et al., 2007) and the respiration volume per time (RVT, Birn et al., 2006) have been applied successfully to denoise BOLD signal from physiological fluctuations after RETROICOR, especially when convolved with a modelled response function, cardiac (Chang et al., 2009) or respiratory (Birn et al., 2008). Various alternatives to RVT to improve breathing-related denoising have been proposed, either to simplify their calculation, that is normally based on the peak detection in the respiratory signal, or to improve the detection of particular changes in the respiratory signal. For instance, Chang and Glover, 2009 proposed to simply use the standard deviation of the respiratory signal, thus avoiding peak detection. Power et al., 2018 suggested to compute the standard deviation of the respiratory envelope in small windows to be more susceptible to breathing changesI think this is not the proper reference by Power for this method. More recently, Harrison et al., 2021 showed that applying an Hilbert transform to compute RVT improves the characterisations of breathing rhythms and the detection of deep breaths. Although the regions impacted by slow variations in cardiac rate and breathing patterns frequently overlap (Chang et al., 2009; Kassinopoulos & Mitsis, 2019), HR and RVT regressor can be used together for better performance (Chang et al., 2009).

Another physiological confound related to RVT consists in spontaneous changes in arterial CO2, called poikilocapnia, which act as a vasoactive process. These fluctuations have been corroborated with Transcranial Doppler ultrasound (TCD) and induce low frequency fluctuations in the BOLD signal. If they are not accounted for, they can induce a bias in the signal estimation in up to a fifth of the cortex (Wise et al., 2004). The pattern of biases induced by poikilocapnia has been found comparable to that of RVT (Chang & Glover, 2009), although accounting for the latter is not suffcient to explain all the variability induced by the former (Golestani et al., 2015). However, the fact that BOLD signal is susceptible to CO2 fluctuations can be conversely seen as an advantage, and used to image cerebral physiology (Moia et al., 2021; Pinto et al., 2021; Zvolanek et al., 2023).

1.4 fMRI Data Analysis

1.4.1 Task fMRI

Conventionally, the analysis of task fMRI data consists of solving a generalized liner model (GLM) analysis. Since researchers have access to the timings of the stimuli, it is possible to model the expected BOLD response to the stimuli and then fit the model to the data. For instance, the onsets of the expected neuronal activity for a given condition can be modeled as an indicator function $p(t)$ (e.g., Dirac functions for event-related designs, or box-car functions for block-designs) convolved with the HRF $h(t)$, sampled at the resolution of the TR (Boynton

et al., 1996; Cohen, 1997; Friston et al., 1998, 2008):

$$x(t) = p \times h(t) \rightarrow x[k] = p \times h(k \cdot \mathrm{TR}). \tag{1.4}$$

Hence, the vector $\mathbf{x} = [x[k]]_{k=1,\ldots,N} \in \mathbb{R}^N$ represents the regressor that models the hypothetical BOLD response for an experimental condition. Then, different regressors either modeling conditions of interest or signals of non interest (e.g., noise-related signals) can be added as the columns of the design matrix $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_L] \in \mathbb{R}^{N \times L}$, where L is the number of regressors, which leads to the well-known GLM equation:

$$\mathbf{y} = \mathbf{X}\beta + \epsilon, \tag{1.5}$$

where the voxel timecourse $\mathbf{y} \in \mathbb{R}^N$ is explained by a linear combination of the regressors in $\mathbf{X}$, weighted by the regression coefficients $\beta \in \mathbb{R}^L$, and the residual error or noise $\epsilon \in \mathbb{R}^N$. The GLM can be solved using ordinary least squares (OLS) under the assumption that the noise is independent and identically distributed (i.i.d.) Gaussian. Hence, the estimates of neuronal activity corresponding to the task conditions are obtained minimizing the residual sum of squares between the fitted model and the measured voxel timecourse. Usually, the number of regressors is much smaller than the number of samples, and thus the GLM is an overdetermined problem. In this case, the solution can be obtained without the need for additional constraints or assumptions (Henson & Friston, 2007). It is important to note that the GLM only provides estimates of the neuronal activity associated with the timings of the modelled stimuli.

1.4.2 Resting-State fMRI

Conversely, fMRI data analysis in unconstrained conditions such as when subjects lying still in the scanner without performing any specific task, i.e. in resting-state, poses a challenge for the application of GLM-based approaches due to the absence of an experimental paradigm that can be used to model the expected BOLD response. Given the increasing popularity of resting-state fMRI in the past decade, various data-driven approaches have emerged to address the analysis of this type of data.

Seed correlation analysis is widely recognized as the predominant approach for examining resting-state fMRI data due to its frequent utilization in the computation of functional connectivity (FC) patterns. It involves measuring the pairwise Pearson correlation between the timecourse of different voxels or regions of the brain, presenting the brain's interregional connections or edges in the form of a FC matrix. Each edge in the matrix represents the strength or intensity of the functional connectivity between two regions. In fact, FC has been extensively utilized to investigate the arrangement of large-scale brain networks (Margulies et al., 2016; Salvador et al., 2005; Yeo et al., 2011), as well as to partition smaller brain structures like the thalamus and striatum (Martino et al., 2008; Tian et al., 2020; Yuan et al., 2015). Notably, FC exhibits subject-specific variations to such an extent that it can serve as a means of individual identification within a diverse population. Numerous studies have demonstrated the presence of distinct

and reliable subject-specific features, commonly referred to as "fingerprints," within the FC matrix (Amico & Goñi, 2018; Finn et al., 2015, 2017; Horien et al., 2019; Jalbrzikowski et al., 2020; Jo et al., 2021; Miranda-Dominguez et al., 2014; Peña-Gómez et al., 2017; Vanderwal et al., 2017; Waller et al., 2017). In fact, these fingerprints have been associated with predictive insights into cognition (Cole et al., 2012; Fong et al., 2019; Rosenberg et al., 2015, 2020; Sripada et al., 2019; Yamashita et al., 2018), personality traits (Adelstein et al., 2011; Dubois et al., 2018; Hsu et al., 2018; Nostro et al., 2018), age (Cabral et al., 2017; Dosenbach et al., 2010; Liem et al., 2017; Nielsen et al., 2018), and disease phenotype (Emerson et al., 2017; Lake et al., 2019; Lynall et al., 2010; Plitt et al., 2015; Svaldi et al., 2021).

Functional connectivity patterns offer a valuable insight into the connectivity among brain regions over a timecourse of data. However, they only provide a single snapshot and do not capture how this connectivity evolves over time, thereby neglecting the dynamics of functional connectivity. (Allen et al., 2012; Di & Biswal, 2020; Hutchison et al., 2013). Dynamic functional connectivity (dFC) in resting-state is commonly investigated using sliding-window approaches (Allen et al., 2012; Hutchison et al., 2013; Lurie et al., 2020; Preti et al., 2017). A sliding-window FC analysis yields a series of time-varying matrices, which are often effectively condensed into a few distinct brain states using clustering techniques (Hutchison et al., 2013). In particular, it has been shown that dFC correlates with underlying neural activity (Keilholz, 2014; Tagliazucchi et al., 2012; Thompson et al., 2013) and behavior (Liégeois et al., 2019). Furthermore, studies have demonstrated that dynamic connectivity fluctuations exhibit lower variability between regions within the same functional networks, while showing higher variability between regions from different networks (Fu et al., 2017). This pattern results in an overall negative correlation with stationary functional connectivity (Thompson & Fransson, 2015; Zhang et al., 2018). However, due to the unconstrained nature of resting-state fMRI, it becomes challenging to ascertain the functional significance of the obtained dynamic connectivity estimates versus their potential derivation from noise (Lindquist et al., 2014). Recently, researchers have also explored dynamic connectivity in the context of participants exposed to complex stimuli, such as movie clips (Di & Biswal, 2020). The utilization of movie stimuli offers the advantage of comparing the time course of dynamic connectivity across participants. If there is a high degree of inter-individual similarity (Hasson et al., 2004; Nastase et al., 2019), it may suggest that the observed dynamics of brain patterns is functionally meaningful and relevant to stimulus processing.

In addition, ongoing developments in the field of dFC methods allow for their operation within individual timeframes. One such method is the use of instantaneous phase synchrony (PS), which provides a reliable measure of connectivity with maximal temporal resolution, comparable to correlation-based methods. Comparing these patterns over time can be achieved by calculating the percentage of time points that exhibit significant phase synchronization throughout the entire scanning duration (Glerean et al., 2012). Alternatively, the leading eigenvectors can be studied for this purpose (Cabral et al., 2017).

Another approach involves employing wavelet transform coherence to explore nonstationary changes in the coupling between fMRI time series. This method calculates coherence and phase lag between two time series as a function of both time and frequency. The selected time series could be either seed timecourses (Chang & Glover, 2010) or timecourses of an ICA (Yaesoubi et al., 2015). Similar to the sliding window approach, temporal dynamics can be identified through the application of clustering algorithms (e.g. k-means) (Cabral et al., 2017; Yaesoubi et al., 2015).

Another alternative method for analyzing single timeframe resting-state fMRI data involves investigating co-activation patterns (CAP) (Chen, Chang, et al., 2015; Liu et al., 2013, 2018; Tagliazucchi et al., 2012). Unlike the phase synchrony approach, CAP analyses focus on identifying simultaneous occurrences of BOLD signal peaks or troughs in different brain regions, disregarding the phase of the signal, assuming that the relationship between the BOLD signal and neural activity is attributed to these brief, transient and sparse co-activation events (Zhang et al., 2020). Typically, CAP analyses employs k-means clustering to group the identified events into distinct CAPs, enabling the identification of temporal dynamics that can potentially be compared to other time-varying resting-state methods such as sliding-window correlation in dFC and PS. Still, the basic approach in CAPs uses the fMRI signal as its input, which is thus subject to the temporal blurring of the hemodynamic response. This phenomenon could lead to the simultaneous co-activation of multiple regions, despite their distinct initial onsets, potentially indicating their association with different components (Rangaprakash et al., 2018). In other words, due to temporal variations in the the hemodynamic response, the BOLD signals of several brain regions might exhibit simultaneous peaks despite the fact that their underlying neuronal activity might have different timings.

To address this potential ambiguity, these neuronal-related events can also be identified by means of hemodynamic deconvolution approaches that remove the blurring effect of the hemodynamic response from the time series (Gaudes et al., 2013; Karahanoğlu et al., 2013; Petridou et al., 2013). Hemodynamic deconvolution is commonly used in investigating psychophysiologic interactions (PPI) within task-based functional connectivity studies (Gerchen et al., 2014; Gitelman et al., 2003) as well as in resting-state fMRI (Di & Biswal, 2015). Recent deconvolution techniques, in contrast to classical PPI analysis, employ sparsity-promoting estimators that assume the dynamics of spontaneous brain activity can be characterized by sparse BOLD events (Gaudes et al., 2013; Karahanoğlu et al., 2013; Petridou et al., 2013; Uruñuela et al., 2023). These approaches are akin to methods using point process analysis to examine sparse BOLD events (Tagliazucchi et al., 2012), and will be the focus of this thesis as introduced in the next section.

Mixing the ideas of hemodynamic deconvolution and co-activation patterns, Karahanoğlu and Van De Ville, 2015 proposed a new approach named innovation-driven co-activation patterns (iCAPs) that is based in transients (i.e. innovations) of the fMRI signal, rather than its peaks. This technique employs the hemodynamic deconvolution algorithm of Total Activation (Karahanoğlu et

al., 2013) to estimate the underlying neural activity prior to applying the CAP analysis, and therefore encode information about transient changes in the signals originating the BOLD timecourses. Evidence from the study conducted by Karahanoğlu and Van De Ville, 2015 using the framework revealed that well-known resting-state networks, including the default mode network, can be subdivided into multiple subsystems with distinct temporal dynamics. This suggests the existence of functionally diverse subnetworks within these networks. Furthermore, by backprojecting iCAPs to deconvolved fMRI volumes, it becomes possible to reconstruct iCAP time courses and assess temporal overlaps between different patterns (Zoller et al., 2019). Notably, it has been observed that, on average, 3 to 4 iCAPs overlap in time, and the associated brain activity persists for 5-10 seconds. This finding may explain the necessity of using a window length of at least 20 seconds to obtain reliable inferences when employing a sliding window approach (Karahanoğlu & Van De Ville, 2015; Preti et al., 2017).

Another way to investigate functional connectivity is by identifying quasi-periodic patterns (QPP) of connectivity, which typically persist for approximately 20 seconds in humans. A repeated-template-averaging algorithm can be employed to detect these patterns in spatiotemporal segments (Majeed et al., 2011). This approach involves iteratively computing sliding window correlations and averaging similar segments of BOLD timepoints until convergence is achieved. As a result, a spatiotemporal averaged template of BOLD dynamics is obtained. Interestingly, these templates often reveal patterns that reflect global signal fluctuations, representing the average time course of the BOLD signal across the entire brain (Bolt et al., 2022; Yousefi et al., 2018). Furthermore, these QPPs have been found to correlate with local infraslow neural activity (Thompson et al., 2014). The predominant QPP typically exhibits a sequence characterized by a transition from strong activation of the default mode network (DMN) and deactivation of sensory and attention networks to the opposite state, with DMN deactivation and activation of sensory and attention networks (Abbas et al., 2019; Yousefi & Keilholz, 2021).

Finally, notice that when the FC strengh between two timecourses (i.e. the edges of the FC matrix) is measured as the pairwise Pearson's correlation, this can be exactly defined in terms of its frame-wise contributions. Therefore, instead of detecting significant instances of FC change from voxel timecourses, these can be identified when multiple timecourses exhibit extreme signal changes simultaneously (i.e. co-activation) (Esfahlani et al., 2020; Faskowitz et al., 2020). These edge-centric FC (eFC) approach has recently gained notable popularity in brain imaging and neuroscience, showing that eFC offers large replicability, stability within individuals across multiple scanning sessions, and reliability across datasets (Faskowitz et al., 2020). Moreover, clustering the eFC has revealed overlapping brain communities that hold promise for studying cognition and behavior beyond the limitations of traditional disjoint brain parcellations. However, the main findings of the edge-centric analyses can be derived from a node-centric FC perspective nFC (i.e., the commonly-used FC matrix) under a static null hypothesis that disregards temporal correlations (Novelli & Razi, 2022). Consequently, the findings obtained with eFC-based

methods can be also applied to nFC-based approaches, such as (i)CAPs or hemodynamic deconvolution, providing a similar set of volumes with large co-activations are used for subsequent analyses.

1.5 Introduction to the Chapters and Aim of the Studies

The aim of this thesis is to develop novel methods for the hemodynamic deconvolution of fMRI data and apply them in resting-state, naturalistic paradigms, and clinical-based studies.

In Chapter 2 the theoretical background of hemodynamic deconvolution is presented with a focus on the two different mathematical formulations, i.e., analysis and synthesis, and the associated algorithms: Total Activation and Paradigm Free Mapping. This chapter will aim to compare the two methods (that work at the individual voxel level), assess whether they are equivalent, and how they compare with co-activation pattern analysis.

Chapter 3 addresses the challenge of determining adequate regularization parameters for hemodynamic deconvolution. The method of stability selection is proposed as a solution to circumvent the need for manual parameter selection. Furthermore, a novel metric based on the area of the stability path is introduced that quantifies the likelihood of a neuronal-related event occurring at a specific voxel and timepoint.

Chapter 4 presents an extension of the hemodynamic deconvolution problem to a multivariate formulation, which enables the identification of neuronal-related events that are shared across multiple voxels. The method is presented and applied to single-echo and multi-echo fMRI data, employs the stability selection method introduced in Chapter 3 for parameter selection, and is validated with a comparison with a generalized linear model as the ground truth. Furthermore, the performance of this novel multivariate method is compared with its univariate counterpart, as well as with another semi-blind multivariate deconvolution technique that has been proposed in the literature, named Hemolearn.

Building upon this multivariate formulation, Chapter 5 introduces a novel method for the simultaneous identification of the neuronal-related events in multiple subjects. The technique is presented as a new tool that allows for the analysis of naturalistic fMRI data at the finest spatial and temporal resolution, unlike other methods that have been proposed in the literature so far. The method is validated in simulated and real data, and its performance is corroborated by correlating the identified neuronal-related events with different features of the movies the subjects watched during the fMRI acquisition.

A novel method, presented in Chapter 6, introduces a paradigm free mapping approach that effectively identifies and decouples global fluctuations in the BOLD signal during the deconvolution process. Compensating for these global events during data preprocessing is challenging, and they can be mistakenly interpreted as neuronally related due to their temporal signature closely resembling the assumed HRF in the deconvolution model. Hence, this decoupling is crucial for optimizing deconvolution approaches, as the estimation accuracy can be

significantly compromised by widespread signal changes caused by head jerks, hardware artifacts, or non-neuronal physiological events (e.g., deep breaths).The method is validated in simulated and real data, and its performance is compared with a GLM analysis as the ground truth.

Finally, Chapter 7 summarizes the main findings of this thesis and outlines future directions for research.

1.6 List of Contributions

Part of the work described in this thesis has been published in peer-reviewed journals and presented in conferences. The following list summarizes the main contributions of this thesis:

1.6.1 Journal and Conference Articles

- Chapter 2: Uruñuela, E., Bolton, T. A., Van De Ville, D., & Caballero-Gaudes, C. (2023). Hemodynamic Deconvolution Demystified: Sparsity-Driven Regularization at Work. Aperture Neuro, vol. 3, 1-25. https://doi.org/10.52294/001c.87574

- Chapter 3: Uruñuela, E., Jones, S., Crawford, A., Shin, W., Oh, S., Lowe, M., & Caballero-Gaudes, C. (2020, July). Stability-Based Sparse Paradigm Free Mapping Algorithm for Deconvolution of Functional MRI Data. In 2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC) (pp. 1092-1095). https://doi.org/10.1109/EMBC44109.2020.9176137

- Chapter 4: Uruñuela, E., Gonzalez-Castillo, J., Zheng, C., Bandettini, P., & Caballero-Gaudes, C. (2023). Whole-Brain Multivariate Hemodynamic Deconvolution for fMRI with Stability Selection. Medical Image Analysis, 103010. https://doi.org/10.1016/j.media.2023.103010

- Chapter 6: Uruñuela, E., Moia, S., & Caballero-Gaudes, C. (2021, April). A Low Rank and Sparse Paradigm Free Mapping Algorithm for Deconvolution of fMRI Data. In 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI) (pp. 1726-1729). https://doi.org/10.1109/ISBI48211.2021.9433821

1.6.2 Conference Abstracts

- Uruñuela, E., Sava-Segal, C., Leung, M., Finn, E.S., Caballero-Gaudes, C. A Multi-Subject Deconvolution Algorithm for the Analysis of Naturalistic fMRI Data. 3rd Annual meeting of the Iberian Chapter of the International Society of Magnetic Resonance in Medicine (2023), Valladolid, Spain.

- Uruñuela, E., Sava-Segal, C., Leung, M., Finn, E.S., Caballero-Gaudes, C. A Multi-Subject Deconvolution Algorithm for the Analysis of Naturalistic fMRI Data. 32nd Annual Meeting of the International Society of Magnetic Resonance in Medicine (2023), Toronto, Canada.
- Uruñuela, E., Moia, S., Caballero-Gaudes, C. A Multi-Echo Low-Rank and Sparse Method to Estimate Neuronal Signal with Less Global Signal Bias. 2nd Annual meeting of the Iberian Chapter of the International Society of Magnetic Resonance in Medicine (2022), Lisbon, Portugal.
- Uruñuela, E., Moia, S., Caballero-Gaudes, C. A Multi-Echo Low-Rank and Sparse Method to Estimate Neuronal Signal with Less Global Signal Bias. 28th Annual Meeting of the Organization for Human Brain Mapping (2022), Glasgow, UK.
- Uruñuela, E., Ferrer, V., Caballero-Gaudes, C. Blind Estimation of Neuronal-Related Activity in fMRI Informed by Co-Fluctuations of Brain Regions. 28th Annual Meeting of the Organization for Human Brain Mapping (2022), Glasgow, UK.
- Uruñuela, E., Moia, S., Caballero-Gaudes, C. A Multi-Echo Low-Rank and Sparse Algorithm That Reduces the Bias of Global Fluctuations on the Estimation of Neuronal Signal. 31st Annual Meeting of the International Society of Magnetic Resonance in Medicine (2022), London, UK.
- Uruñuela, E., Moia, S., Caballero-Gaudes, C. Synthesis-Based Paradigm Free Mapping and Analysis-Based Total Activation Operate Identically. 27th Annual Meeting of the Organization for Human Brain Mapping (2021), Virtual.
- Uruñuela, E., Moia, S., Caballero-Gaudes, C. Simultaneous Blind Estimation of Global Fluctuations and Neuronal-Related Activity from fMRI Data. 1st Annual meeting of the Iberian Chapter of the International Society of Magnetic Resonance in Medicine (2021), Virtual.
- Uruñuela, E., Moia, S., Caballero-Gaudes, C. Synthesis-Based Paradigm Free Mapping and Analysis-Based Total Activation Are Identical For Temporal Deconvolution of fMRI Data. 1st Annual meeting of the Iberian Chapter of the International Society of Magnetic Resonance in Medicine (2021), Virtual.
- Uruñuela, E., Moia, S., Caballero-Gaudes, C. fMRI deconvolution with Synthesis-Based Paradigm Free Mapping and Analysis-Based Total Activation Operate Identically. 30th Annual Meeting of the International Society of Magnetic Resonance in Medicine (2021), Virtual.
- Uruñuela, E., Moia, S., Caballero-Gaudes, C. Low-Rank and Sparse Simultaneous Blind Estimation of Global Fluctuations and Neuronal-Related Activity from fMRI Data. 30th Annual Meeting of the International Society of Magnetic Resonance in Medicine (2021), Virtual.

- Uruñuela, E., Jones, S., Crawford, A., Shin, W., Oh, S., Lowe, M., Caballero-Gaudes, C. Improving Deconvolution of fMRI Signal with Sparse Paradigm Free Mapping Using Stability Selection. 26th Annual Meeting of the Organization for Human Brain Mapping (2020), Virtual.
- Uruñuela, E., Moia, S., Gonzalez-Castillo, J., Caballero-Gaudes, C. Deconvolution of Multi-Echo Functional MRI Data with Multivariate Multi-Echo Sparse Paradigm Free Mapping. 27th Annual Meeting of the International Society of Magnetic Resonance in Medicine (2019), Montreal, Canada.

Chapter

2

Synthesis-Based and Analysis-Based Hemodynamic Deconvolution for fMRI

This chapter was published as Uruñuela, E., Bolton, T. A., Van De Ville, D., & Caballero-Gaudes, C. (2023). Hemodynamic deconvolution demystified: Sparsity-driven regularization at work. Aperture Neuro, vol. 3, 1–25. DOI: https://doi.org/10.52294/001c.87574.

Deconvolution of the hemodynamic response is an important step to access short timescales of brain activity recorded by functional magnetic resonance imaging (fMRI). Albeit conventional deconvolution algorithms have been around for a long time (e.g., Wiener deconvolution), recent state-of-the-art methods based on sparsity-pursuing regularization are attracting increasing interest to investigate brain dynamics and connectivity with fMRI. This chapter describes the principles of hemodynamic deconvolution in fMRI, and revisits the main concepts underlying two main sparsity-promoting deconvolution algorithms, Paradigm Free Mapping and Total Activation, in an accessible manner. Despite their apparent differences in the formulation, these methods are shown theoretically equivalent as they represent the synthesis and analysis sides of the same problem, respectively. This chapter demonstrates this equivalence in practice with their best-available implementations using both simulations, with different signal-to-noise ratios, and experimental fMRI data acquired during a motor task and resting-state. The chapter also evaluates the parameter settings that lead to equivalent results, and showcases the operation of these algorithms in comparison with co-activation pattern analysis. This note is useful for practitioners interested in gaining a better understanding of state-of-the-art hemodynamic deconvolution, and aims to answer questions regarding the differences between the two methods.

2.1 Introduction

Functional magnetic resonance imaging (fMRI) data analysis is often directed to identify and disentangle the neural processes that occur in different brain regions during task or at rest. As the blood oxygenation level-dependent (BOLD) signal of fMRI (Ogawa et al., 1990) is only a proxy for neuronal activity mediated through neurovascular coupling (Logothetis et al., 2001), an intermediate step that estimates the activity-inducing signal, at the timescale of fMRI, from the BOLD timeseries can be useful. Conventional analysis of task fMRI data relies on the general linear models (GLM) to establish statistical parametric maps of brain activity by regression of the empirical timecourses against hypothetical ones built from the knowledge of the experimental paradigm (Boynton et al., 1996; Cohen, 1997; Friston et al., 1998, 2008). However, timing information of the paradigm can be unknown, inaccurate, or insufficient in some scenarios such as naturalistic stimuli, resting-state, or clinically-relevant assessments.

Deconvolution and methods alike are aiming to estimate neuronal activity by undoing the blurring effect of the hemodynamic response, characterized as a hemodynamic response function (HRF)[1]. Given the inherently ill-posed nature of hemodynamic deconvolution, due to the sluggish characteristics of the HRF, the key is to introduce additional constraints in the estimation problem that are typically expressed as regularization terms. For instance, the so-called Wiener deconvolution is expressing a "minimal energy" constraint on the deconvolved signal, and has been used in the framework of psychophysiological interaction analysis to compute the interaction between a seed's activity-inducing timecourse and an experimental modulation (Di & Biswal, 2018; Freitas et al., 2020; Gerchen et al., 2014; Gitelman et al., 2003; Glover, 1999). Complementarily, the interest in deconvolution has increased to explore time-varying activity in resting-state fMRI data (Bolton, Morgenroth, et al., 2020; Keilholz et al., 2017; Lurie et al., 2020; Preti et al., 2017). In that case, the aim is to gain better insights of the neural signals that drive functional connectivity at short time scales, as well as learning about the spatio-temporal structure of functional components that dynamically construct resting-state networks and their interactions (Karahanoğlu & Van De Ville, 2017).

Deconvolution of the resting-state fMRI signal has illustrated the significance of transient, sparse spontaneous events (Allan et al., 2015; Petridou et al., 2013) that refine the hierarchical clusterization of functional networks (Karahanoğlu et al., 2013) and reveal their temporal overlap based on their signal innovations not only in the human brain (Karahanoğlu & Van De Ville, 2015), but also in the spinal cord (Kinany et al., 2020). Similar to task-related studies, deconvolution allows to investigate modulatory interactions within and between resting-state functional networks (Di & Biswal, 2013, 2015). In addition, decoding of the deconvolved spontaneous events allows to decipher the flow of spontaneous thoughts and actions across different cognitive and sensory domains while at

[1] Note that the term deconvolution is also alternatively employed to refer to the estimation of the hemodynamic response shape assuming a known activity-inducing signal or neuronal activity (Casanova et al., 2008; Ciuciu et al., 2003; Goutte et al., 2000; Marrelec et al., 2002).

rest (Gonzalez-Castillo et al., 2019; Karahanoğlu & Van De Ville, 2015; Tan et al., 2017). Beyond findings on healthy subjects, deconvolution techniques have also proven its utility in clinical conditions to characterize functional alterations of patients with a progressive stage of multiple sclerosis at rest (Bommarito et al., 2021), to find functional signatures of prodromal psychotic symptoms and anxiety at rest on patients suffering from schizophrenia (Zoeller et al., 2019), to detect the foci of interictal events in epilepsy patients without an EEG recording (Karahanoglu et al., 2013; Lopes et al., 2012), or to study functional dissociations observed during non-rapid eye movement sleep that are associated with reduced consolidation of information and impaired consciousness (Tarun et al., 2021).

The algorithms for hemodynamic deconvolution can be classified based on the assumed hemodynamic model and the optimization problem used to estimate the neuronal-related signal. Most approaches assume a linear time-invariant model for the hemodynamic response that is inverted by means of variational (regularized) least squares estimators (Caballero-Gaudes et al., 2019; Cherkaoui et al., 2019; Costantini et al., 2022; Gaudes et al., 2012; Gaudes et al., 2010, 2013; Gitelman et al., 2003; Glover, 1999; Hernandez-Garcia & Ulfarsson, 2011; Hütel et al., 2021; Karahanoğlu et al., 2013), logistic functions (Bush & Cisler, 2013; Bush, Zhou, et al., 2015; Loula et al., 2018), probabilistic mixture models (Pidnebesna et al., 2019), convolutional autoencoders (Hütel et al., 2018) or nonparametric homomorphic filtering (Sreenivasan et al., 2015). Alternatively, several methods have also been proposed to invert non-linear models of the neuronal and hemodynamic coupling (Aslan et al., 2016; Friston et al., 2008; Havlicek et al., 2011; Madi & Karameh, 2017; Penny et al., 2005; Riera et al., 2004; Ruiz-Euler et al., 2018).

Among the variety of approaches, those based on regularized least squares estimators have been employed more often due to their appropriate performance at small spatial scales (e.g., voxelwise). Relevant for this work, two different formulations can be established for the regularized least-squares deconvolution problem, either based on a synthesis- or analysis-based model (Elad et al., 2007; Ortelli & van de Geer, 2019). On the one hand, Paradigm Free Mapping is based on a synthesis formulation that is solved by means of regularized least-squares estimators such as ridge-regression (Gaudes et al., 2010) or LASSO (Gaudes et al., 2013). The rationale of the synthesis-based model is that we know or suspect that the true signal (here, the neuronally-driven BOLD component of the fMRI signal) can be represented as a linear combination of predefined patterns or dictionary atoms (for instance, the hemodynamic response function). On the other hand, Total Activation is based on a analysis formulation that is solved with a regularized least-squares estimator using generalized total variation (Karahanoglu et al., 2011; Karahanoğlu et al., 2013). The rationale of the analysis-based approach considers that the true signal is analyzed by some relevant hemodynamic operator (Khalidov et al., 2011) and the resulting signal is sparse in time.

The users of these algorithms have often questioned about the similarities and differences between the two methods and which one is better. To clarify this point, this chapter initially presents the theory behind these two deconvolution

approaches: Paradigm Free Mapping (PFM) (Gaudes et al., 2013) — available in AFNI as 3dPFM[2] and 3dMEPFM[3] for single-echo and multi-echo data, respectively — and Total Activation (TA) (Karahanoğlu et al., 2013) — available as part of the iCAPs toolbox[4]. The chapter describes the similarities and differences in their analytical formulations, and how they can be related to each other. Next, their performance is assessed controlling for a fair comparison on simulated and experimental data. Finally, the chapter discusses their benefits and shortcomings and conclude with our vision on potential extensions and developments.

2.2 Theory

2.2.1 Notations and Definitions

Matrices of size N rows and M columns are denoted by boldface capital letters, e.g., $\mathbf{X} \in \mathbb{R}^{N \times M}$, whereas column vectors of length N are denoted as boldface lowercase letters, e.g., $\mathbf{x} \in \mathbb{R}^N$. Scalars are denoted by lowercase letters, e.g., k. Continuous functions are denoted by brackets, e.g., $h(t)$, while discrete functions are denoted by square brackets, e.g., $x[k]$. The Euclidean norm of a matrix $\mathbf{X}$ is denoted as $\|\mathbf{X}\|_2$, the ℓ_1-norm is denoted by $\|\mathbf{X}\|_1$ and the Frobenius norm is denoted by $\|\mathbf{X}\|_F$. The discrete integration ($\mathbf{L}$) and difference ($\mathbf{D}$) operators are defined as:

$$\mathbf{L} = \begin{bmatrix} 1 & 0 & \dots & & \\ 1 & 1 & 0 & \dots & \\ 1 & 1 & 1 & 0 & \dots \\ \vdots & \ddots & \ddots & \ddots & \ddots \end{bmatrix}, \quad \mathbf{D} = \begin{bmatrix} 1 & 0 & \dots & & \\ 1 & -1 & 0 & \dots & \\ 0 & \ddots & \ddots & \ddots & \dots \\ \vdots & \ddots & 0 & 1 & -1 \end{bmatrix}.$$

2.2.2 General Linear Model Analysis

A conventional general linear model (GLM) analysis puts forward a number of regressors incorporating the knowledge about the paradigm or behavior. For instance, the timing of epochs for a certain condition can be modeled as an indicator function $p(t)$ (e.g., Dirac functions for event-related designs or box-car functions for block-designs) convolved with the hemodynamic response function (HRF) $h(t)$, and sampled at the repetition time (TR) resolution (Boynton et al., 1996; Cohen, 1997; Friston, Jezzard, & Turner, 1994; Friston et al., 1998):

$$x(t) = p * h(t) \rightarrow x[k] = p * h(k \cdot \mathrm{TR}).$$

The vector $\mathbf{x} = [x[k]]_{k=1,\dots,N} \in \mathbb{R}^N$ then constitutes the regressor modelling the hypothetical response, and several of them can be stacked as columns of the

[2] https://afni.nimh.nih.gov/pub/dist/doc/program_help/3dPFM.html
[3] https://afni.nimh.nih.gov/pub/dist/doc/program_help/3dMEPFM.html
[4] https://c4science.ch/source/iCAPs/

design matrix $\mathbf{X} = [\mathbf{x}_1 \ldots \mathbf{x}_L] \in \mathbb{R}^{N \times L}$, leading to the well-known GLM formulation:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}, \tag{2.1}$$

where the empirical timecourse $\mathbf{y} \in \mathbb{R}^N$ is explained by a linear combination of the regressors in $\mathbf{X}$ weighted by the parameters in $\boldsymbol{\beta} \in \mathbb{R}^L$ and corrupted by additive noise $\mathbf{e} \in \mathbb{R}^N$. Under independent and identically distributed Gaussian assumptions of the latter, the maximum likelihood estimate of the parameter weights reverts to the ordinary least-squares estimator; i.e., minimizing the residual sum of squares between the fitted model and measurements. The number of regressors L is typically much less than the number of measurements N, and thus the regression problem is over-determined and does not require additional constraints or assumptions (Henson & Friston, 2007).

In the deconvolution approach, no prior knowledge of the hypothetical response is taken into account, and the purpose is to estimate the deconvolved activity-inducing signal $\mathbf{s}$ from the measurements $\mathbf{y}$, which can be formulated as the following signal model

$$\mathbf{y} = \mathbf{H}\mathbf{s} + \mathbf{e}, \tag{2.2}$$

where $\mathbf{H} \in \mathbb{R}^{N \times N}$ is a Toeplitz matrix that represents the discrete convolution with the HRF, and $\mathbf{s} \in \mathbb{R}^N$ is a length-N vector with the unknown activity-inducing signal. Note that the temporal resolution of the activity-inducing signal and the corresponding Toeplitz matrix is generally assumed to be equal to the TR of the acquisition, but it could also be higher if an upsampled estimate is desired. Despite the apparent similarity with the GLM equation, there are two important differences. First, the multiplication with the design matrix of the GLM is an expansion as a weighted linear combination of its columns, while the multiplication with the HRF matrix represents a convolution operator. Second, determining $\mathbf{s}$ is an ill-posed problem given the nature of the HRF. As it can be seen intuitively, the convolution matrix $\mathbf{H}$ is highly collinear (i.e., its columns are highly correlated) due to large overlap between shifted HRFs (see Figure 2.2C), thus introducing uncertainty in the estimates of $\mathbf{s}$ when noise is present. Consequently, additional assumptions under the form of regularization terms (or priors) in the estimate are needed to reduce their variance. In the least squares sense, the optimization problem to solve is given by

$$\hat{\mathbf{s}} = \arg\min_{\mathbf{s}} \frac{1}{2}\|\mathbf{y} - \mathbf{H}\mathbf{s}\|_2^2 + \Omega(\mathbf{s}). \tag{2.3}$$

The first term quantifies data fitness, which can be justified as the log-likelihood term derived from Gaussian noise assumptions, while the second term $\Omega(\mathbf{s})$ brings in regularization and can be interpreted as a prior on the activity-inducing signal. For example, the ℓ_2-norm of $\mathbf{s}$ (i.e., $\Omega(\mathbf{s}) = \lambda\|\mathbf{s}\|_2^2$) is imposed for ridge regression or Wiener deconvolution, which introduces a trade-off between the data fit term and "energy" of the estimates that is controlled by the regularization parameter λ. Other well-known regularized terms are related to the elastic net (i.e., $\Omega(\mathbf{x}) = \lambda_1\|\mathbf{x}\|_2^2 + \lambda_2\|\mathbf{x}\|_1$) (Zou & Hastie, 2005).

2.2.3 Paradigm Free Mapping

In Paradigm Free Mapping (PFM), the formulation of Equation (2.3) was considered equivalently as fitting the measurements using the atoms of the HRF dictionary (i.e., columns of $\mathbf{H}$) with corresponding weights (entries of $\mathbf{s}$). This model corresponds to a synthesis formulation. In Gaudes et al., 2013 a sparsity-pursuing regularization term was introduced on $\mathbf{s}$, which in a strict way reverts to choosing $\Omega(\mathbf{s}) = \lambda\|\mathbf{s}\|_0$ as the regularization term and solving the optimization problem (Bruckstein et al., 2009). However, finding the optimal solution to the problem demands an exhaustive search across all possible combinations of the columns of $\mathbf{H}$. Hence, a pragmatic solution is to solve the convex-relaxed optimization problem for the l_1-norm, commonly known as Basis Pursuit Denoising (Chen et al., 2001) or equivalently as the least absolute shrinkage and selection operator (LASSO) (Tibshirani, 1996):

$$\hat{\mathbf{s}} = \arg\min_{\mathbf{s}} \frac{1}{2}\|\mathbf{y} - \mathbf{H}\mathbf{s}\|_2^2 + \lambda\|\mathbf{s}\|_1, \tag{2.4}$$

which provides fast convergence to a global solution. Imposing sparsity on the activity-inducing signal implies that it is assumed to be well represented by a reduced subset of few non-zero coefficients at the fMRI timescale, which in turn trigger event-related BOLD responses. Hereinafter, this assumption is referred to as the spike model. However, even if PFM was developed as a spike model, its formulation in Equation (2.4) can be extended to estimate the innovation signal, i.e., the derivative of the activity-inducing signal, as shown in Section 2.2.5.

2.2.4 Total Activation

Alternatively, deconvolution can be formulated as if the signal to be recovered directly fits the measurements and at the same time satisfies some suitable regularization, which leads to

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \frac{1}{2}\|\mathbf{y} - \mathbf{x}\|_2^2 + \Omega(\mathbf{x}). \tag{2.5}$$

Under this analysis formulation, total variation (TV), i.e., the ℓ_1-norm of the derivative $\Omega(\mathbf{x}) = \lambda\|\mathbf{D}\mathbf{x}\|_1$, is a powerful regularizer since it favors recovery of piecewise-constant signals (Chambolle, 2004). Going beyond, the approach of generalized TV introduces an additional differential operator $\mathbf{D_H}$ in the regularizer that can be tailored as the inverse operator of a linear system (Karahanoglu et al., 2011), that is, $\Omega(\mathbf{x}) = \lambda\|\mathbf{D}\mathbf{D_H}\mathbf{x}\|_1$. In the context of hemodynamic deconvolution, Total Activation is proposed for which the discrete operator $\mathbf{D_H}$ is derived from the inverse of the continuous-domain linearized Balloon-Windkessel model (Buxton et al., 1998; Friston et al., 2000). The interested reader is referred to (Karahanoglu et al., 2011; Karahanoğlu et al., 2013; Khalidov et al., 2011) for a detailed description of this derivation.

Therefore, the solution of the Total Activation (TA) problem

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \frac{1}{2}\|\mathbf{y} - \mathbf{x}\|_2^2 + \lambda\|\mathbf{D}\mathbf{D_H}\mathbf{x}\|_1 \tag{2.6}$$

will yield the activity-related signal $\mathbf{x}$ for which the activity-inducing signal $\mathbf{s} = \mathbf{D_H x}$ and the so-called innovation signal $\mathbf{u} = \mathbf{Ds}$, i.e., the derivate of the activity-inducing signal, will also be available, as they are required for the regularization. Modeling the activity-inducing signal based on the innovation signal is referred to as the block model. Nevertheless, even if TA was originally developed as a block model, its formulation in Equation (2.6) can be made equivalent to the spike model as shown in Section 2.2.5.

2.2.5 Unifying Both Perspectives

PFM and TA are based on the synthesis- and analysis-based formulation of the deconvolution problem, respectively. They are also tailored for the spike and block model, respectively. In the first case, the recovered deconvolved signal is synthesized to be matched to the measurements, while in the second case, the recovered signal is directly matched to the measurements but needs to satisfy its analysis in terms of deconvolution. This also corresponds to using the forward or backward model of the hemodynamic system, respectively. Hence, it is possible to make both approaches equivalent (Elad et al., 2007)[5].

To start with, TA can be made equivalent to PFM by adapting it for the spike model; i.e., when removing the derivative operator $\mathbf{D}$ of the regularizer in Equation (2.6), it can be readily verified that replacing in that case $\mathbf{x} = \mathbf{Hs}$ leads to identical equations and thus both assume a spike model, since $\mathbf{H}$ and $\mathbf{D_H}$ will cancel out each other (Karahanoglu et al., 2011)[6].

Conversely, the PFM spike model can also accommodate the TA block model by modifying Equation (2.4) with the forward model $\mathbf{y} = \mathbf{HLu} + \mathbf{e}$. Here, the activity-inducing signal $\mathbf{s}$ is rewritten in terms of the innovation signal $\mathbf{u}$ as $\mathbf{s} = \mathbf{Lu}$ where the matrix $\mathbf{L}$ is the first-order integration operator (Cherkaoui et al., 2019; Uruñuela et al., 2020). This way, PFM can estimate the innovation signal $\mathbf{u}$ as follows:

$$\hat{\mathbf{u}} = \arg\min_{\mathbf{u}} \frac{1}{2} \|\mathbf{y} - \mathbf{HLu}\|_2^2 + \lambda \|\mathbf{u}\|_1, \tag{2.7}$$

and becomes equivalent to TA by replacing $\mathbf{u} = \mathbf{DD_H x}$, and thus adopting the block model. Based on the previous equations (Equation (2.4)), (Equation (2.6)) and (Equation (2.7)), it is clear that both PFM and TA can operate under the spike and block models, providing a convenient signal model according to the different assumptions of the underlying neuronal-related signal. This work evaluates the core of the two techniques; i.e., the regularized least-squares problem with temporal regularization without considering the spatial regularization term originally incorporated in TA. For the remainder of this paper, the PFM and TA formalisms with both spike and block models are used.

[5]Without dwelling into technicalities, for total variation, this equivalence is correct up to the constant, which is in the null space of the derivative operator.

[6]Again, this holds up to elements of the null space.

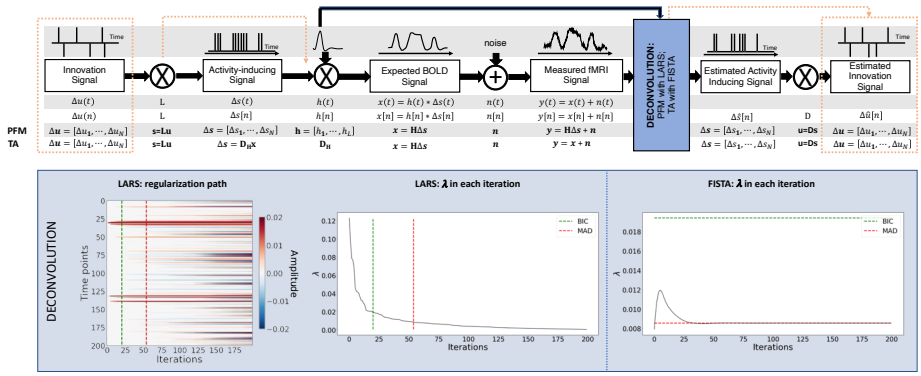

Figure 2.1: Flowchart detailing the different steps of the fMRI signal and the deconvolution methods described. The orange arrows indicate the flow to estimate the innovation signals, i.e., the derivative of the activity-inducing signal. The blue box depicts the iterative modus operandi of the two algorithms used in this paper to solve the paradigm free mapping (PFM) and Total Activation (TA) deconvolution problems. The plot on the left shows the regularization path obtained with the least angle regression (LARS) algorithm, where the x-axis illustrates the different iterations of the algorithm, the y-axis represents points in time, and the color describes the amplitude of the estimated signal. The middle plot depicts the decreasing values of λ for each iteration of LARS as the regularization path is computed. The green and red dashed lines in both plots illustrate the Bayesian information criterion (BIC) and median absolute deviation (MAD) solutions, respectively. Comparatively, the changes in λ when the fast iterative shrinkage-thresholding algorithm (FISTA) method is made to converge to the MAD estimate of the noise are shown on the right. Likewise, the λ corresponding to the BIC and MAD solutions are shown with dashed lines.

2.2.6 Algorithms and Parameter Selection

Despite their apparent resemblance, the practical implementations of the PFM and TA methods proposed different algorithms to solve the corresponding optimization problem and select an adequate regularization parameter λ (Gaudes et al., 2013; Karahanoğlu et al., 2013). The PFM implementation available in AFNI employs the least angle regression (LARS) (Efron et al., 2004), whereas the TA implementation uses the fast iterative shrinkage-thresholding algorithm (FISTA) (Beck & Teboulle, 2009). The blue box in Figure 2.1 provides a descriptive view of the iterative modus operandi of the two algorithms.

On the one hand, LARS is a homotopy approach that computes all the possible solutions to the optimization problem and their corresponding value of λ; i.e., the regularization path, and the solution according to the Bayesian Information Criterion (BIC) (Schwarz, 1978), was recommended as the most appropriate in the case of PFM approaches since Akaike Information Criterion

(AIC) often tends to overfit the signal (Caballero-Gaudes et al., 2019; Gaudes et al., 2013).

On the other hand, FISTA is an extension of the classical gradient algorithm that provides fast convergence for large-scale problems. In the case of FISTA though, the regularization parameter λ must be selected prior to solving the problem, but can be updated in every iteration so that the residuals of the data fit converge to an estimated noise level of the data $\hat{\sigma}$:

$$\lambda^{n+1} = \frac{N\hat{\sigma}}{\frac{1}{2}\|\mathbf{y}-\mathbf{x}^n\|_F^2}\lambda^n, \tag{2.8}$$

where x^n is the n^{th} iteration estimate, λ^n and λ^{n+1} are the n^{th} and $n+1^{th}$ iteration values for the regularization parameter λ, and N is the number of points in the time-course. The pre-estimated noise level can be obtained as the median absolute deviation (MAD) of the fine-scale wavelet coefficients (Daubechies, order 3) of the fMRI timecourse. The MAD criterion has been adopted in TA (Karahanoğlu et al., 2013). Of note, similar formulations based on the MAD estimate have also been applied in PFM formulations (Gaudes et al., 2012; Gaudes et al., 2011).

2.3 Methods

2.3.1 Simulated data

In order to compare the two methods while controlling for their correct performance, a simulation scenario was created, which can be found in the GitHub repository shared in Section 2.6. For the sake of illustration, here the simulations correspond to a timecourse with a duration of 400 seconds (TR = 2 s) where the activity-inducing signal includes 5 events, which are convolved with the canonical HRF. Different noise sources (physiological, thermal, and motion-related) were also added and three different scenarios with varying signal-to-noise ratios (SNR = 20, 10, 3 dB) were simulated, which represent high, medium and low contrast-to-noise ratios as shown in Figure 2.2A. Noise was created following the procedure in (Gaudes et al., 2013) as the sum of uncorrelated Gaussian noise and sinusoidal signals to simulate a realistic noise model with thermal noise, cardiac and respiratory physiological fluctuations, respectively. The physiological signals were generated as

$$\sum_{i=1}^{2} \frac{1}{2^{i-1}} \left(\sin\left(2\pi f_{r,i}\, t + \phi_{r,i}\right) + \sin\left(2\pi f_{c,i}\, t + \phi_{c,i}\right)\right), \tag{2.9}$$

with up to second-order harmonics per cardiac ($f_{c,i}$) and respiratory ($f_{r,i}$) component that were randomly generated following normal distributions with variance 0.04 and mean $i f_r$ and $i f_c$, for $i = [1,2]$. The fundamental frequencies were set to $f_r = 0.3$ Hz for the respiratory component (Birn et al., 2006) and $f_c = 1.1$ Hz for the cardiac component (Shmueli et al., 2007). The phases of each

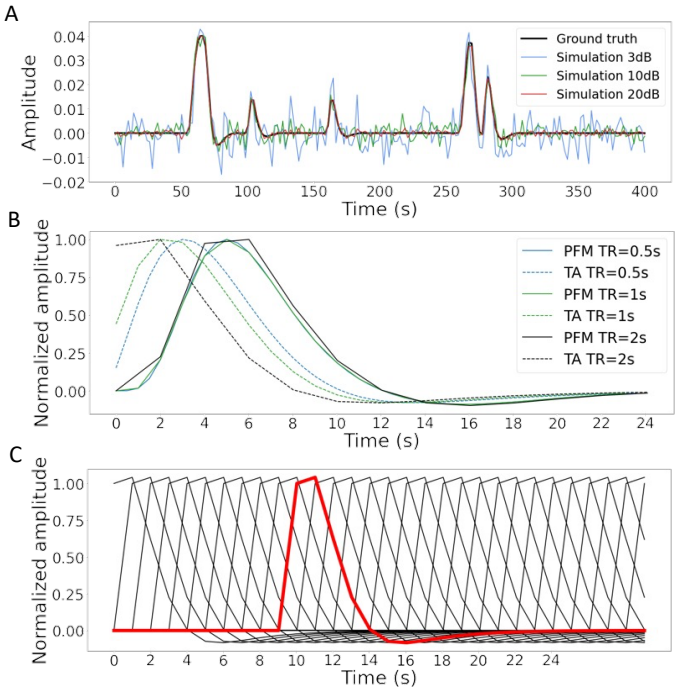

Figure 2.2: A) Simulated signal with different SNRs (20 dB, 10 dB and 3 dB) and ground truth given in signal percentage change (SPC). B) Canonical HRF models typically used by PFM (solid line) and TA (dashed line) at TR = 0.5 s (blue), TR = 1 s (green) and TR = 2 s (black). Without loss of generality, the waveforms are scaled to unit amplitude for visualization. C) Representation of shifted HRFs at TR = 2 s that build the design matrix for PFM when the HRF model has been matched to that in TA. The red line corresponds to one of the columns of the HRF matrix.

harmonic ϕ were randomly selected from a uniform distribution between 0 and 2π radians. To simulate physiological noise that is proportional to the change in BOLD signal, a variable ratio between the physiological (σ_P) and the thermal (σ_0) noise was modeled as $\sigma_P/\sigma_0 = a(tSNR)^b + c$, where $a = 5.01 \times 10^{-6}$, $b = 2.81$, and $c = 0.397$, following the experimental measures available in Table 3 from (Triantafyllou et al., 2005)).

2.3.2 Experimental Data

To compare the performance of the two approaches as well as illustrate their operation, two representative experimental datasets were employed.

Motor task dataset: One healthy subject was scanned in a 3T MR scanner

(Siemens) under a Basque Center on Cognition, Brain and Language Review Board-approved protocol. T2*-weighted multi-echo fMRI data was acquired with a simultaneous-multislice multi-echo gradient echo-planar imaging sequence, kindly provided by the Center of Magnetic Resonance Research (University of Minnesota, USA) (Feinberg et al., 2010; Moeller et al., 2010; Setsompop et al., 2011), with the following parameters: 340 time frames, 52 slices, Partial-Fourier = 6/8, voxel size = $2.4 \times 2.4 \times 3$ mm^3, TR = 1.5 s, TEs = 10.6/28.69/46.78/64.87/82.96 ms, flip angle = 70^o, multiband factor = 4, GRAPPA = 2.

During the fMRI acquisition, the subject performed a motor task consisting of five different movements (left-hand finger tapping, right-hand finger tapping, moving the left toes, moving the right toes and moving the tongue) that were visually cued through a mirror located on the head coil. These conditions were randomly intermixed every 16 seconds, and were only repeated once the entire set of stimuli were presented. Data preprocessing consisted of first, discarding the first 10 volumes of the functional data to achieve a steady state of magnetization. Then, image realignment to the skull-stripped single-band reference image (SBRef) was computed on the first echo, and the estimated rigid-body spatial transformation was applied to all other echoes (Jenkinson & Smith, 2001; Jenkinson et al., 2012). A brain mask obtained from the SBRef volume was applied to all the echoes and the different echo timeseries were optimally combined (OC) voxelwise by weighting each timeseries contribution by its T2* value (Posse et al., 1999). AFNI (Cox, 1996) was employed for a detrending of up to 4^{th}-order Legendre polynomials, within-brain spatial smoothing (3 mm FWHM) and voxelwise signal normalization to percentage change. Finally, distortion field correction was performed on the OC volume with Topup (Andersson et al., 2003), using the pair of spin-echo EPI images with reversed phase encoding acquired before the ME-EPI acquisition (Glasser et al., 2016).

Resting-state datasets: One healthy subject was scanned in a 3T MR scanner (Siemens) under a Basque Center on Cognition, Brain and Language Review Board-approved protocol. Two runs of T2*-weighted fMRI data were acquired during resting-state, each with 10 min duration, with 1) a standard gradient-echo echo-planar imaging sequence (monoband) (TR = 2000 ms, TE = 29 ms, flip-angle = 78^o, matrix size = 64×64, voxel size = $3 \times 3 \times 3$ mm^3, 33 axial slices with interleaved acquisition, slice gap = 0.6 mm) and 2) a simultaneous-multislice gradient-echo echo-planar imaging sequence (multiband factor = 3, TR = 800 ms, TE = 29 ms, flip-angle = 60^o, matrix size = 64×64, voxel size = $3 \times 3 \times 3$ mm^3, 42 axial slices with interleaved acquisition, no slice gap). Single-band reference images were also collected in both resting-state acquisitions for head motion realignment. Field maps were also obtained to correct for field distortions.

During both acquisitions, participants were instructed to keep their eyes open, fixating a white cross that they saw through a mirror located on the head coil, and not to think about anything specific. The data was pre-processed using AFNI (Cox, 1996). First, volumes corresponding to the initial 10 seconds were removed to allow for a steady-state magnetization. Then, the voxel time-series

were despiked to reduce large-amplitude deviations and slice-time corrected. Inhomogeneities caused by magnetic susceptibility were corrected with FUGUE (FSL) using the field map images (Jenkinson et al., 2012). Next, functional images were realigned to a base volume (monoband: volume with the lowest head motion; multiband: single-band reference image). Finally, a simultaneous nuisance regression step was performed comprising up to 6^{th}-order Legendre polynomials, low-pass filtering with a cutoff frequency of 0.25 Hz (only on multiband data to match the frequency content of the monoband), 6 realignment parameters plus temporal derivatives, 5 principal components of white matter (WM), 5 principal components of lateral ventricle voxels (anatomical CompCor) (Behzadi et al., 2007) and 5 principal components of the brain's edge voxels (Patriat et al., 2015). WM, cerebrospinal fluid (CSF) and brain's edge-voxel masks were obtained from Freesurfer tissue and brain segmentations. In addition, scans with potential artifacts were identified and censored when the euclidean norm of the temporal derivative of the realignment parameters (ENORM) was larger than 0.4, and the proportion of voxels adjusted in the despiking step exceeded 10%.

2.3.3 Selection of the Hemodynamic Response Function

In their original formulations, PFM and TA specify the discrete-time HRF in different ways. For PFM, the continuous-domain specification of the canonical double-gamma HRF (Henson & Friston, 2007) is sampled at the TR and then put as shifted impulse responses to build the matrix $\mathbf{H}$. In the case of TA, however, the continuous-domain linearized version of the balloon-windkessel model is discretized to build the linear differential operator in $\mathbf{D_H}$. While the TR only changes the resolution of the HRF shape for PFM, the impact of an equivalent impulse response of the discretized differential operator at different TR is more pronounced. As shown in Figure 2.2B, longer TR leads to equivalent impulse responses of TA that are shifted in time, provoking a lack of the initial baseline and rise of the response. The reader is referred to Figure A.1 to see the differences in the estimation of the activity-inducing and innovation signals when both methods use the HRF in their original formulation. To avoid differences between PFM and TA based on their built-in HRF, the synthesis operator $\mathbf{H}$ was built with shifted versions of the HRF given by the TA analysis operator (e.g., see Figure 2.2C for the TR=2s case).

2.3.4 Selection of the Regularization Parameter

The simulated data was used to compare the performance of the two deconvolution algorithms with both BIC and MAD criteria to set the regularization parameter λ (see Section 2.2.6). Here, the evaluation also included investigating if the algorithms behave differently in terms of the estimation of the activity-inducing signal $\hat{\mathbf{s}}$ using the spike model described in Equation (2.4) and the block model based on the innovation signal $\hat{\mathbf{u}}$ in Equation (2.7).

For selection based on the BIC, LARS was initally performed with the PFM deconvolution model to obtain the solution for every possible λ in the regularization path. Then, the values of λ corresponding to the BIC optimum were adopted to solve the TA deconvolution model by means of FISTA.

For a selection based on the MAD estimate of the noise, the temporal regularization in its original form for TA was applied, whereas for PFM the selected λ corresponds to the solution whose residuals have the closest standard deviation to the estimated noise level of the data $\hat{\sigma}$.

2.3.5 Analyses in Experimental fMRI Data

Difference between approaches: To assess the discrepancies between both approaches when applied on experimental fMRI data, the square root of the sum of squares of the differences (RSSD) between the activity-inducing signals estimated with PFM and TA were calculated on the three experimental datasets as

$$\mathrm{RSSD} = \sqrt{\frac{1}{N}\sum_{k=1}^{N}(\hat{s}_{\mathrm{PFM}}[k] - \hat{s}_{\mathrm{TA}}[k])^2}, \tag{2.10}$$

where N is the number of timepoints of the acquisition. The RSSD of the innovation signals $\hat{\mathbf{u}}$ was computed equally.

Task fMRI data: In the analysis of the motor task data, the performance of PFM and TA was evaluated in comparison with a conventional General Linear Model analysis (3dDeconvolve in AFNI) that takes advantage of the information about the duration and onsets of the motor trials. Given the block design of the motor task, this comparison is only made with the block model.

Resting-state fMRI data: The usefulness of deconvolution approaches in the analysis of resting-state data where information about the timings of neuronal-related BOLD activity cannot be predicted is also illustrated. Apart from being able to explore individual maps of deconvolved activity (i.e., innovation signals, activity-inducing signals, or hemodynamic signals) at the temporal resolution of the acquisition (or deconvolution), here the average extreme points of the activity-inducing and innovation maps (given that these examples do not have a sufficient number of scans to perform a clustering step) is calculated to illustrate how popular approaches like co-activation patterns (CAPs) (Liu et al., 2018; Tagliazucchi et al., 2012) and innovation-driven co-activation patterns (iCAPs) (Karahanoğlu & Van De Ville, 2015) can be applied on the deconvolved signals to reveal patterns of coordinated brain activity. To achieve this, the average time-series was calculated in a seed of 9 voxels located in the precuneus, supramarginal gyrus, and occipital gyri independently, and solve the deconvolution problem to find the activity-inducing and innovation signals in the seeds. Then, a 95^{th} percentile threshold was applied and the maps of the time-frames that survive the threshold were averaged. Finally, the same procedure was applied to the original— i.e., non-deconvolved— signal in the seed and compare the results with the widely-used seed correlation approach.

2.4 Results

2.4.1 Performance Based on the Regularization Parameter

Figure 2.3A shows the regularization paths of PFM and TA side by side obtained for the spike model of Equation (2.4) for SNR=3 dB. The solutions for all three SNR conditions are shown in Figures A.2 and A.3. Starting from the maximum λ corresponding to a null estimate and for decreasing values of λ, LARS computes a new estimate at the value of λ that reduces the sparsity promoted by the l_1-norm and causes a change in the active set of non-zero coefficients of the estimate (i.e., a zero coefficient becomes non-zero or vice versa) as shown in the horizontal axis of the heatmaps. Vertical dashed lines depict the selection of the regularization parameter based on the BIC, and thus, the colored coefficients indicated by these depict the estimated activity-inducing signal $\hat{\mathbf{s}}$. Figure 2.3B illustrates the resulting estimates of the activity-inducing and activity-related hemodynamic signals when basing the selection of λ on the BIC for SNR=3 dB. Given that the regularization paths of both approaches are identical, it can be clearly observed that the BIC-based estimates are identical too for the corresponding λ. Thus, Figure 2.3A, Figure 2.3B, Figure A.2 and Figure A.3 demonstrate that, regardless of the simulated SNR condition, the spike model of both deconvolution algorithms produces identical regularization paths when the same HRF and regularization parameters are applied, and hence, identical estimates of the activity-inducing signal $\hat{\mathbf{s}}$ and neuronal-related hemodynamic signal $\hat{\mathbf{x}}$.

Likewise, Figure 2.3C demonstrates that the regularization paths for the block model defined in Equations (2.6) and (2.7) also yield virtually identical estimates of the innovation signals for both PFM and TA methods. Again, the BIC-based selection of λ is identical for both PFM and TA. As illustrated in Figure 2.3D, the estimates of the innovation signal $\mathbf{u}$ also show no distinguishable differences between the algorithms. Figure 2.3A-D demonstrate that both PFM and TA yield equivalent regularization paths and estimates of the innovation signal and activity-inducing signal regardless of the simulated SNR condition when applying the same HRF and regularization parameters with the block and spike models.

As for selecting λ with the MAD criterion defined in Equation (2.8), Figure 2.3E depicts the estimated activity-inducing and activity-related signals for the simulated low-SNR setting using the spike model, while Figure 2.3F shows the estimated signals corresponding to the block model. Both plots in Figure 2.3E and F depict nearly identical results between PFM and TA with both models. Given that the regularization paths of both techniques are identical, minor dissimilarities are owing to the slight differences in the selection of λ due to the quantization of the values returned by LARS.

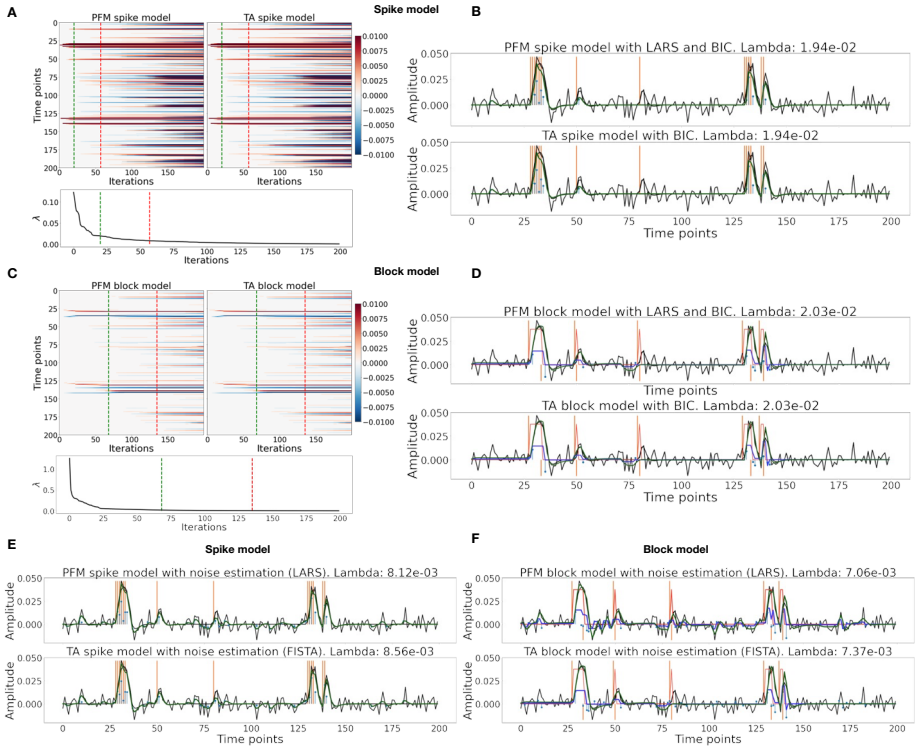

Figure 2.3: (A) Heatmap of the regularization paths of the activity-inducing signals (spike model) estimated with PFM and TA as a function of λ for the simulated data with SNR = 3 dB (x-axis: increasing number of iterations or λ as given by LARS; y-axis: time; color: amplitude). Vertical lines denote iterations corresponding to the BIC (dashed line) and MAD (dotted line) selection of λ. (B) Estimated activity-inducing (blue) and activity-related (green) signals with a selection of λ based on the BIC. Orange and red lines depict the ground truth. (C) Heatmap of the regularization paths of the innovation signals (block model) estimated with PFM and TA as a function of λ for the simulated data with SNR = 3 dB. (D) Estimated innovation (blue), activity-inducing (darker blue), and activity-related (green) signals with a selection of λ based on the BIC. (E) Activity-inducing and activity-related (fit, $\mathbf{x}$) signals estimated with PFM (top) and TA (bottom) when λ is selected based on the MAD method with the spike model, and (F) with the block model for the simulated data with SNR = 3 dB.

2.4.2 Performance on Experimental Data

Figure 2.4 depicts the RSSD maps revealing differences between PFM and TA estimates for the spike (Figure 2.4A and C) and block (Figure 2.4B and D) models when applied to the three experimental fMRI datasets. The RSSD values

are virtually negligible (i.e., depicted in yellow) in most of the within-brain voxels and lower than the amplitude of the estimates of the activity-inducing and innovation signals. Based on the maximum value of the range shown in each image, it can be observed that the similarity between both approaches is more evident for the spike model (with both selection criteria) and the block model with the BIC selection. However, given the different approaches used for the selection of the regularization parameter λ based on the MAD estimate of the noise (i.e., converging so that the residuals of FISTA are equal to the MAD estimate of the noise for TA vs. finding the LARS residual that is closest to the MAD estimate of the noise), higher RSSD values can be observed with the largest differences occurring in gray matter voxels. These areas also correspond to low values of λ (see Figure A.4) and MAD estimates of the noise (see Figure A.5), while the highest values are visible in regions with signal droupouts, ventricles, and white matter. These differences that arise from the different approaches to find the optimal regularization parameter based on the MAD estimate of the noise can be clearly seen in the root sum of squares (RSS) of the estimates (Figure A.6). These differences are also observable in the ATS calculated from estimates obtained with the MAD selection as shown in Figure A.9. However, the identical regularization paths shown in Figure A.7 demonstrate that both methods perform equivalently on experimental data (see estimates of innovation signal obtained with an identical selection of λ in Figure A.8). Hence, the higher RSSD values originate from the different methods to find the optimal regularization parameter based on the MAD estimate of the noise that yield different solutions as shown by the dashed vertical lines in Figure A.7.

Figure 2.5 depicts the results of the analysis of the Motor dataset with the PFM and TA algorithms using the BIC selection of λ (see Figure A.9 for results with MAD selection), as well as a conventional GLM approach. The Activation Time Series (top left), calculated as the sum of squares of all voxel amplitudes (positive vs. negative) for a given moment in time, obtained with PFM and TA show nearly identical patterns. These ATS help to summarize the four dimensional information available in the results across the spatial domain and identify instances of significant BOLD activity. The second to sixth rows show the voxel timeseries and the corresponding activity-related, activity-inducing and innovation signals obtained with PFM using the BIC criterion of representative voxels in the regions activated in each of the motor tasks. The TA-estimated time-series are not shown because they were virtually identical. The maps shown on the right correspond to statistical parametric maps obtained with the GLM for each motor condition ($p < 0.001$) as well as the maps of the PFM and TA estimates at the onsets of individual motor events (indicated with arrows in the timecourses). The estimated activity-related, activity-inducing and innovation signals clearly reveal the activity patterns of each condition in the task, as they exhibit a BOLD response locked to the onset and duration of the conditions. Overall, activity maps of the innovation signal obtained with PFM and TA highly resemble those obtained with a GLM for individual events, with small differences arising from the distinct specificity of the GLM and deconvolution analyses. Notice that the differences observed with the different approaches to

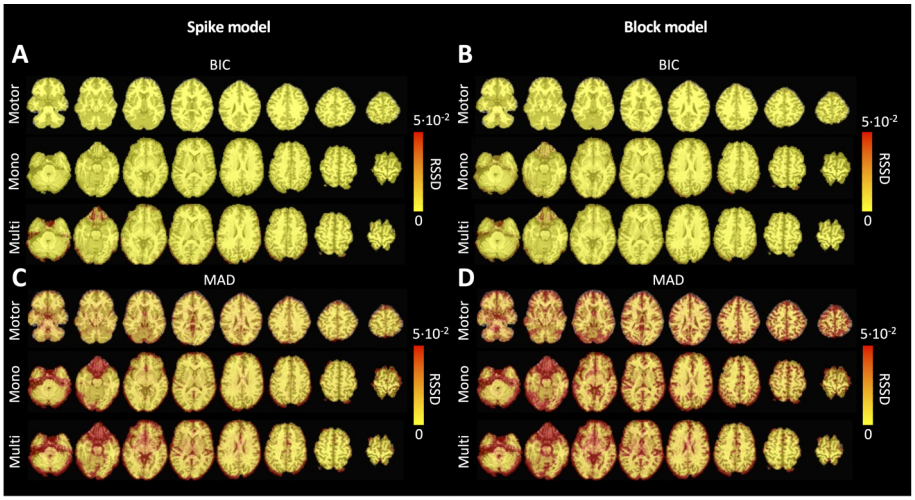

Figure 2.4: Square root of the sum of squared differences (RSSD) between the estimates obtained with PFM and TA for (A) spike model (activity-inducing signal) and BIC selection of λ, (B) block model (innovation signal) and BIC selection, (C) spike model (activity-inducing signal) and MAD selection, (D) block model (innovation signal) and MAD selection. RSSD maps are shown for the three experimental fMRI datasets: the motor task (Motor), the monoband resting-state (Mono), and the multiband resting-state (Multi) datasets.

select λ based on the MAD estimate shown in Figure 2.4 are reflected on the ATS shown in Figure A.9 as well.

As an illustration of the insights that deconvolution methods can provide in the analysis of resting-state data, Figure 2.6 depicts the average activity-inducing and innovation maps of common resting-state networks obtained from thresholding and averaging the activity-inducing and innovation signals, respectively, estimated from the resting-state multiband data using PFM with a selection of λ based on the BIC. The average activity-inducing maps obtained via deconvolution show spatial patterns of the default mode network (DMN), dorsal attention network (DAN), and visual network (VIS) that highly resemble the maps obtained with conventional seed correlation analysis using Pearson's correlation, and the average maps of extreme points of the signal (i.e., with no deconvolution). With deconvolution, the average activity-inducing maps seem to depict more accurate spatial delineation (i.e., less smoothness) than those obtained from the original data, while maintaining the structure of the networks. The BIC-informed selection of λ yields spatial patterns of average activity-inducing and innovation maps that are more sparse than those obtained with a selection of λ based on the MAD estimate (see Figure A.10). Furthermore, the spatial patterns of the average innovation maps based on the innovation signals

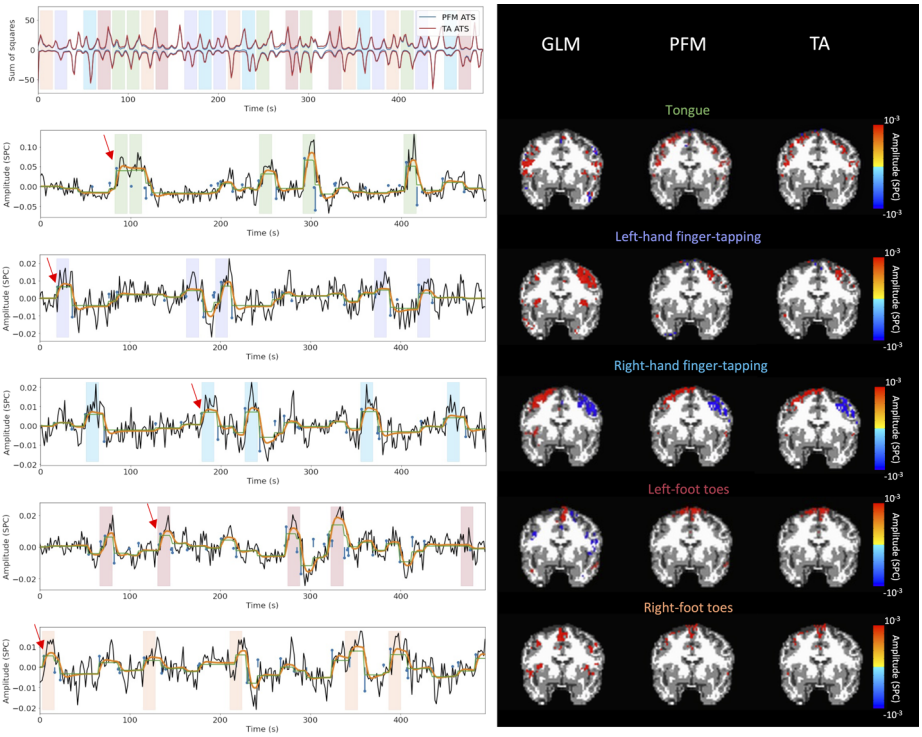

Figure 2.5: Activity maps of the motor task using a selection of λ based on the BIC estimate. Row 1: Activation time-series (ATS) of the innovation signals estimated by PFM (in blue) or TA (in red) calculated as the sum of squares of all voxels at every timepoint. Positive-valued and negative-valued contributions were separated into two distinct time-courses. Color-bands indicate the onset and duration of each condition in the task (green: tongue motion, purple: left-hand finger-tapping, blue: right-hand finger-tapping, red: left-foot toes motion, orange: right-foot toes motion). Rows 2-6: time-series of a representative voxel for each task with the PFM-estimated innovation (blue), PFM-estimated activity-inducing (green), and activity-related (i.e., fitted, orange) signals, with their corresponding GLM, PFM, and TA maps on the right (representative voxels indicated with green arrows). Amplitudes are given in signal percentage change (SPC). The maps shown on the right are sampled at the time-points labeled with the red arrows and display the innovation signals at these moments across the whole brain.

using the block model yield complementary information to those obtained with the activity-inducing signal since iCAPs allow to reveal regions with synchronous innovations, i.e., with the same upregulating and downregulating events. For instance, it is interesting to observe that the structure of the visual network

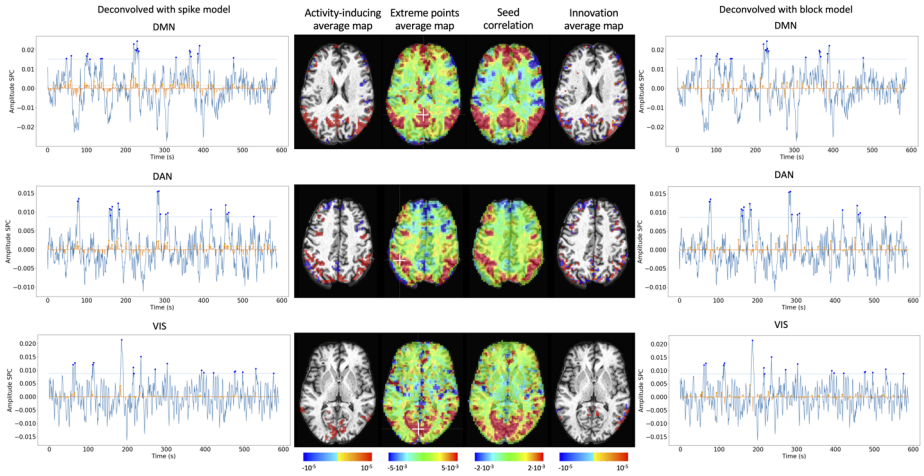

Figure 2.6: Average activity-inducing (left) and innovation (right) maps obtained from PFM-estimated activity-inducing and innovation signals, respectively, using a BIC-based selection of λ. Time-points selected with a 95^{th} percentile threshold (horizontal lines) are shown over the average time-series (blue) in the seed region (white cross) and the deconvolved signals, i.e., activity inducing (left) and innovation (right) signals (orange). Average maps of extreme points and seed correlation maps are illustrated in the center.

nearly disappears in its corresponding average innovation maps, suggesting the existence of different temporal neuronal patterns across voxels in the primary and secondary visual cortices.

2.5 Discussion and Conclusion

Hemodynamic deconvolution can be formulated using a synthesis- and analysis-based approach as proposed by PFM and TA, respectively. This work demonstrates that the theoretical equivalence of both approaches is confirmed in practice given virtually identical results when the same HRF model and equivalent regularization parameters are employed. Hence, it can be argued that previously observed differences in performance can be explained by specific settings, such as the HRF model and selection of the regularization parameter (as shown in Figures 2.4, A.6 and A.7), convergence thresholds, as well as the addition of a spatial regularization term in the spatiotemporal TA formulation (Karahanoğlu et al., 2013). For instance, the use of PFM with the spike model in Tan et al., 2017 was seen not to be adequate due to the prolonged trials in the paradigm, which better fit the block model as described in Equation (2.7). However, given the equivalence of the temporal deconvolution, incorporating extra spatial or temporal regularization terms in the optimization problem

should not modify this equivalence providing convex operators are employed. For a convex optimization problem, with a unique global solution, iterative shrinkage thresholding procedures alternating between the different regularization terms guarantee convergence, such as the generalized forward-backward splitting (Raguet et al., 2013) algorithm originally employed for TA.

Our findings are also in line with the equivalence of analysis and synthesis methods in under-determined cases ($N \leq V$) demonstrated in (Elad et al., 2007) and (Ortelli & van de Geer, 2019). Still, this chapter has shown that a slight difference in the selection of the regularization parameter can lead to small differences in the estimated signals when employing the block model with the BIC selection of λ. However, since their regularization paths are equivalent, the algorithms can easily be forced to converge to the same selection of λ, thus resulting in identical estimated signals.

Nevertheless, the different formulations of analysis and synthesis deconvolution models bring along different kinds of flexibility. One notable advantage of PFM is that it can readily incorporate any HRF as part of the synthesis operator (Elad et al., 2007), only requiring the sampled HRF at the desired temporal resolution, which is typically equal to the TR of the acquisition. Conversely, TA relies upon the specification of the discrete differential operator that inverts the HRF, which needs to be derived either by the inverse solution of the sampled HRF impulse response, or by discretizing a continuous-domain differential operator motivated by a biophysical model. The more versatile structure of PFM allows for instance an elegant extension of the algorithm to multi-echo fMRI data (Caballero-Gaudes et al., 2019) where multiple measurements relate to a common underlying signal. Therefore, the one-to-many synthesis scenario (i.e., from activity-inducing to several activity-related signals) is more cumbersome to express using TA. In other words, a set of differential operators should be defined and the differences between their outputs constrained. In contrast, the one-to-many analysis scenario (i.e., from the measurements to several regularizing signals) is more convenient to be expressed by TA, for example combining spike and block regularizers. While the specification of the differential operator in TA only indirectly controls the HRF, the use of the derivative operator to enforce the block model, instead of the integrator in PFM, impacts positively the stability and rate of the convergence of the optimization algorithms. Moreover, analysis formulations can be more suitable for online applications that are still to be explored in fMRI data, but are employed for calcium imaging deconvolution (Friedrich et al., 2017; Jewell et al., 2019), and which have been applied for offline calcium deconvolution (Farouj et al., 2020).

Moreover, deconvolution techniques can be used before more downstream analysis of brain activity in terms of functional network organization as they estimate interactions between voxels or brain regions that occur at the activity-inducing level, and are thus less affected by the slowness of the hemodynamic response compared to when the BOLD signals are analyzed directly. In particular, hemodynamic deconvolution approaches hold a close parallelism to recent methodologies aiming to understand the dynamics of neuronal activations and interactions at short temporal resolution and that focus on extreme events of the

fMRI signal. As an illustration, Figure 2.6 shows that the innovation- or activity-inducing CAPs computed from deconvolved events in a single resting-state fMRI dataset closely resemble the conventional CAPs computed directly from extreme events of the fMRI signal (Cifre, Zarepour, et al., 2020; Cifre, Flores, et al., 2020; Liu & Duyn, 2013; Liu et al., 2013, 2018; Rolls et al., 2021; Tagliazucchi et al., 2011, 2012, 2016; Zhang et al., 2020). Similarly, it can be hypothesized that these extreme events will also show a close resemblance to intrinsic ignition events (Deco & Kringelbach, 2017; Deco et al., 2017). As shown in the maps, deconvolution approaches can offer a more straightforward interpretability of the activation events and resulting functional connectivity patterns. Here, CAPs were computed as the average of spatial maps corresponding to the events of a single dataset. Beyond simple averaging, clustering algorithms (e.g., K-means and consensus clustering) can be employed to discern multiple CAPs or iCAPs at the whole-brain level for a large number of subjects. Previous findings based on iCAPs have for instance revealed organizational principles of brain function during rest (Karahanoğlu & Van De Ville, 2015) and sleep (Tarun et al., 2021) in healthy controls, next to alterations in 22q11ds (Zoeller et al., 2019) and multiple sclerosis (Bommarito et al., 2021).

Next to CAPs-inspired approaches, dynamic functional connectivity has recently been investigated with the use of co-fluctuations and edge-centric techniques (Esfahlani et al., 2020; Faskowitz et al., 2020; Jo et al., 2021; Sporns et al., 2021; van Oort et al., 2018). The activation time series shown in Figure 2.5A aims to provide equivalent information to the root of sum of squares timecourses used in edge-centric approaches, where timecourses with peaks delineate instances of significant brain activity. Future work could address which type of information is redundant or distinct across these frameworks. These examples illustrate that deconvolution techniques can be employed prior to other computational approaches and could serve as an effective way of denoising the fMRI data. Hence, an increase in the number of studies that take advantage of the potential benefits of using deconvolution methods prior to functional connectivity analyses can be expected.

Even though the two approaches examined here provide alternative representations of the BOLD signals in terms of innovation and activity-inducing signals, their current implementations have certain limitations and call for further developments or more elaborate models, where some of them have been partially addressed in the literature. One relevant focus is to account for the variability in HRF that can be observed in different regions of the brain. First, variability in the temporal characteristics of the HRF can arise from differences in stimulus intensity and patterns, as well as with short inter-event intervals like in fast cognitive processes or experimental designs (Chen et al., 2021; Polimeni & Lewis, 2021; Sadaghiani et al., 2009; Yeşilyurt et al., 2008). Similarly, the HRF shape at rest might differ from the canonical HRF commonly used for task-based fMRI data analysis. A wide variety of HRF patterns could be elicited across the whole brain and possible detected with sufficiently large signal-to-noise ratio. For instance, Gonzalez-Castillo et al., 2012 showed two gamma-shaped responses at the onset and the end of the evoked trial, respectively. This unique HRF shape

would be deconvolved as two separate events with the conventional deconvolution techniques. The impact of HRF variability could be reduced using structured regularization terms along with multiple basis functions (Gaudes et al., 2012) or procedures that estimate the HRF shape in an adaptive fashion in both analysis (Farouj et al., 2019) and synthesis formulations (Cherkaoui et al., 2021).

Another avenue of research consists in leveraging spatial information by adopting multivariate deconvolution approaches that operate at the whole-brain level, instead of working voxelwise and beyond regional regularization terms (e.g., as proposed in Karahanoğlu et al., 2013). Operating at the whole-brain level would open the way for methods that consider shared neuronal activity using mixed norm regularization terms (Uruñuela-Tremiño et al., 2019), as described in Chapter 4, or can capture long-range neuronal cofluctuations using low rank decompositions (Cherkaoui et al., 2021). For example, multivariate deconvolution approaches could yield better localized activity patterns while reducing the effect of global fluctuations such as respiratory artifacts, which cannot be modelled at the voxel level with a multivariate sparse and low-rank model (Uruñuela et al., 2021), as described in Chapter 6.

Similar to solving other inverse problems by means of regularized estimators, the selection of the regularization parameter is critical to correctly estimate the neuronal-related signal. Hence, methods that take advantage of a more robust selection of the regularization parameter could considerably yield more reliable estimates of the neuronal-related signal. For instance, the stability selection procedure (Meinshausen & Bühlmann, 2010; Uruñuela et al., 2020)could be included to the deconvolution problem to ensure that the estimated coefficients are obtained with high probability. This approach is described in Chapter 3 (Uruñuela et al., 2020). Furthermore, an important issue of regularized estimation is that the estimates are biased with respect to the true value. In that sense, the use of non-convex $\ell_{p,q}$-norm regularization terms (e.g., $p < 1$) can reduce this bias while maintaining the sparsity constraint, at the cost of potentially converging to a local minima of the regularized estimation problem. In practice, these approaches could avoid the optional debiasing step that overcomes the shrinkage of the estimates and obtain a more accurate and less biased fit of the fMRI signal (Caballero-Gaudes et al., 2019; Gaudes et al., 2013). Finally, cutting-edge developments on physics-informed deep learning techniques for inverse problems (Akçakaya et al., 2021; Cherkaoui et al., 2020; Monga et al., 2021; Ongie et al., 2020) could be transferred for deconvolution by considering the biophysical model of the hemodynamic system and could potentially offer algorithms with reduced computational time and more flexibility.

2.6 Code and Data Availability

The code and materials used in this work can be found in the following GitHub repository: https://github.com/eurunuela/pfm_vs_ta. The reader can explore different simulation parameters (e.g., SNR, varying HRF options and mismatch between algorithms, TR, number of events, onsets, and durations) in the

provided Jupyter notebooks. Similarly, the experimental data can be found in https://osf.io/f3ryg/.

Chapter

3

Stability-Based Sparse Paradigm Free Mapping

This chapter was published as Uruñuela, E., Jones, S., Crawford, A., Shin, W., Oh, S., Lowe, M., & Caballero-Gaudes, C. (2020). Stability-based sparse paradigm free mapping algorithm for deconvolution of functional MRI data. 2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC). DOI: 10.1109/EMBC44109.2020.9176137.

This chapter proposes the implementation of a subsampling approach based on stability selection that avoids the choice of any regularization parameter for hemodynamic deconvolution with sparity-promoting regularized least squares estimators. The proposed method is implemented to operate with the Sparse Paradigm Free Mapping (SPFM) algorithm and its performance is evaluated on real fMRI data and compared with both the original SPFM algorithm, which used model selection criteria to select the parameters, and a conventional analysis with a general linear model (GLM) that is aware of the temporal model of the neuronal-related activity. The results demonstrate that SPFM with stability selection yields activation maps with higher resemblance to the maps obtained with GLM analyses and offers improved detection of neuronal-related events over SPFM, particularly in scenarios with low contrast-to-noise ratio.

3.1 Introduction

In the preceding chapter, deconvolution approaches were discussed in the context of functional magnetic resonance imaging (fMRI) data analysis as they offer the remarkable ability to estimate neuronal-related activity without the need for prior information on the timings of the blood oxygenation level-dependent (BOLD)

events via a linear time-invariant model (i.e., a forward model of the BOLD response) that is then inverted by means of regularized least-squares estimators to deconvolve the neuronal-related activity at each voxel (Gaudes et al., 2010, 2013; Gitelman et al., 2003; Hernandez-Garcia & Ulfarsson, 2011; Karahanoğlu et al., 2013; Khalidov et al., 2011). In particular, the sparse Paradigm Free Mapping (SPFM) method (Gaudes et al., 2013), which is the basis of this work, employs sparsity-promoting regularization terms based on the L1-norm of the estimates (e.g., using the LASSO or the Dantzig Selector). Importantly, inverse problem solving is linked to a dilemma that has yet to be solved: the selection of the regularization parameters that yield accurate estimates. As described in the previous chapter, methods based on model selection criteria after the computation of the entire regularization path (Gaudes et al., 2013) or iterative procedures so that the variance of the residuals after deconvolution is equal to a prior estimate of the noise variance (Karahanoğlu et al., 2013) have been previously used in the literature for parameter tuning due to their reduced computational cost. Yet, these methods offer no information about the appropriateness of the selected parameters.

This chapter proposes the use of the subsampling approach of stability selection (Meinshausen & Bühlmann, 2010) with the SPFM algorithm (Gaudes et al., 2013) to avoid the choice of any regularization parameter and account for the likelihood of the different possible estimates in the regularization path. Although stability selection has been previously proposed in fMRI data analysis, for example in the estimation of functional connectivity matrices from partial correlations with sparse estimators (Ryali et al., 2012) and to detect change points in time-varying functional connectivity with the graphical lasso (Cribben et al., 2013), its application for the deconvolution of the fMRI signal is innovative. Further, this chapter implements a novel procedure that enables to benefit from the computational speed of the least angle regression algorithm (Efron et al., 2004) in combination with the robustness of stability selection.

This chapter uses a modification of the original SPFM formulation called block model –and introduced in Chapter 2– that computes estimates of the innovation signal of the neuronal-related signal (i.e., defining its changes) (Cherkaoui et al., 2019; Karahanoğlu et al., 2013; Uruñuela et al., 2023), rather than the signal itself. The block model formulation fits the data used in this study better as it improves the estimation of neuronal-related events with long, sustained activity (Cherkaoui et al., 2019; Karahanoğlu et al., 2013; Uruñuela et al., 2023) that cannot be adequately described by conventional spike-like models (Gaudes et al., 2010, 2013; Khalidov et al., 2011). Nevertheless, the proposed stability selection procedure can be readily implemented for the spike model. The chapter is organized as follows: in Section 3.2 the signal model and the stability-based SFPM algorithm are introduced; in Section 3.3, the results of applying this new algorithm on experimental fMRI data are presented and compare them to the previous SPFM algorithm.

3.2 Signal Model and Deconvolution with Stability-Based Paradigm Free Mapping

For the sake of completeness, this section revises the signal model for hemodynamic deconvolution, which was already described in Chapter 2. In fMRI data analysis, the signal of a voxel $y(t)$ is commonly modelled as the convolution of an underlying neuronal-related signal $s(t)$ with the hemodynamic response function (HRF) $h(t)$, plus a white noise component: $y(t) = h(t) * s(t) + n(t)$, or $\mathbf{y} = \mathbf{Hs} + \mathbf{n}$ in a discrete-time matrix notation. Typically, the neuronal-related signal $s(t)$ is represented as a train of Dirac impulses at the fMRI timescale associated with the experimental design. This model of the neuronal-related signal has been adopted by previous deconvolution algorithms (Gaudes et al., 2010, 2013; Hernandez-Garcia & Ulfarsson, 2011) relying on regularized least-squares estimators as follows:

$$\hat{\mathbf{s}} = \arg\min_{\mathbf{s}} \frac{1}{2}\|\mathbf{y} - \mathbf{Hs}\|_2^2 + \lambda\|\mathbf{s}\|_p, \tag{3.1}$$

where the L_p-norm $\|\mathbf{s}\|_p$ penalizes the amplitude of the coefficients of the neuronal-related signal, e.g., $p = 2$ (i.e., ridge regression) and $p = 1$ (i.e., LASSO) were employed in (Gaudes et al., 2010) and (Gaudes et al., 2013), respectively. Instead of the on/off pattern described by Dirac impulses, the neuronal-related signal $\mathbf{s}$ can also be represented as a piecewise constant signal in terms of its innovation signal $\mathbf{u}$ (i.e., its first derivative in time). Defining $\mathbf{s} = \mathbf{Lu}$ where $\mathbf{L}$ corresponds to the discrete integration operator (Cherkaoui et al., 2019), the signal model can be written as:

$$\mathbf{y} = \mathbf{HLu} + \mathbf{n}, \tag{3.2}$$

where $\mathbf{y}, \mathbf{u}, \mathbf{s}, \mathbf{n} \in \mathbb{R}^N$, $L \in \mathbb{R}^{NxN}$, and $\mathbf{H} \in \mathbb{R}^{NxN}$ is the Toeplitz convolution matrix with shifted HRFs, where N is the number of observations of the fMRI signal. The signal $\mathbf{u}$ will represent those instances when significant changes in the neuronal-related activity occur. Since the innovation signal $\mathbf{u}$ is sparser than the neuronal-related signal $\mathbf{s}$, it is also a more adequate representation if the temporal deconvolution of the fMRI time series of each voxel is performed with L1-norm regularized estimators as follows:

$$\hat{\mathbf{u}} = \arg\min_{\mathbf{u}} \frac{1}{2}\|\mathbf{y} - \mathbf{HLu}\|_2^2 + \lambda\|\mathbf{u}\|_1. \tag{3.3}$$

3.2.1 Combining Stability Selection with Least Angle Regression

An appropriate choice of the regularization parameter λ in Equation (3.1) or Equation (3.3) is crucial for appropriate hemodynamic deconvolution. Several techniques to select it have already been proposed, such as based on the Bayesian Information Criterion (Gaudes et al., 2013). However, these techniques do not provide a solution that is robust regardless of the different characteristics the

data may show (e.g., signal-to-noise ratio, occurrence and duration of neuronal events).

This problem can be overcome by implementing a novel procedure based on the stability selection approach (Meinshausen & Bühlmann, 2010). This procedure generates $T = 100$ surrogate datasets $y_i (i = 1, \ldots, T)$ where the original voxel time series is randomly subsampled to retain 60% of its time points. Then, the optimization problem in Equation (3.3) is solved for each surrogate dataset, where the model matrix $\mathbf{H}$ is subsampled accordingly. Then, the stability paths of the signal $\mathbf{u}$ for each surrogate i and each time point t (i.e., u_t^i) are computed, which represent the probability of the coefficient being non-zero for a given λ. Originally, the stability selection approach operates by solving Equation (3.3) for a predefined set of λ values, for example by means of the fast iterative shrinkage thresholding algorithm (FISTA) (Beck & Teboulle, 2009). Instead, this chapter proposes to use the least angle regression (LARS) algorithm (Efron et al., 2004), which computes the entire regularization path for an optimal decreasing set of λ values and is faster than FISTA (Beck & Teboulle, 2009) for our purposes. Then, for each surrogate, the estimate $u_(\lambda_l, t)^i$ at the regularization parameter λ_l and time point t is binarized as $c_(\lambda_l, t)^i = 0$ if $u_(\lambda_l, t)^i = 0$ or $c_(\lambda_l, t)^i = 1$ otherwise. To overcome the fact that solving Equation (3.3) with the LARS algorithm will generate a different set of λ values in each subsampled surrogate, a new set of λ values is created. This new set contains all of the regularization parameters from all of the surrogate-specific regularization paths in decreasing order. The coefficients $c_(\lambda_l, t)^i$ remain 0 or 1 according to the preceding value of λ_l^i corresponding to the surrogate-specific regularization path computed by LARS. This step allows us to calculate the probabilities that construct the stability paths as the ratio of surrogates where each coefficient $u_(\lambda, t)$ is different from 0 at each λ.

Furthermore, unlike in the original stability selection procedure that sets a given probability threshold to select the final set of non-zero coefficients, this implementation calculates the area under the curve (AUC) of the stability paths of each coefficient u_t as follows:

$$\mathrm{AUC}_t = \frac{\sum_{l=1}^{L} \lambda_l P\left(u_{\lambda_l, t} \neq 0\right)}{\sum_{l=1}^{L} \lambda_l}, \tag{3.4}$$

where

$$P\left(u_{\lambda_l, t} \neq 0\right) = T^{-1} \sum_{i=1}^{T} c_{\lambda_l, t}^i, \tag{3.5}$$

represents the selection probability of coefficient $u_(\lambda, t)$ for a particular choice of the regularization parameter λ_l, and L is the total number of regularization parameters from all of the LARS regularization paths. Hence, the voxelwise time series AUC(t) reveals the most prominent coefficients indicating the probability of activation at each time point.

3.2.2 Thresholding and Debiasing

Afterwards, the AUC time series for each voxel are thresholded to identify those instances with high probability of a significant change in neuronal-related activity occurring. This threshold is based on a given percentile (or maximum) of the AUC values in a region of interest where no BOLD signal changes related to neuronal activity are assumed to occur (or can be detected). For example, the threshold can be set to the 99th percentile of the AUC values of deep white matter (DWM) voxels (see Section 3.4).

Finally, it is recommended to remove the bias in the estimates of the neuronal-related signal owing to the L1-norm regularization term. For the signal model in Equation (3.1) used in the original SPFM approach (Gaudes et al., 2013), a debiased estimate of $\mathbf{s}$ can be obtained by solving a least squares problem with a selection of non-zero AUC coefficients.

Rather, in the signal model with the innovation signal, the selected non-zero coefficients of $\mathbf{u}$ are used to define a matrix $\mathbf{A}$ whose columns are activation segments with piecewise constant unit between two non-zero coefficients of $\mathbf{u}$ (Zoller et al., 2019). A final debiased estimate of $\mathbf{s}$ is obtained by solving the following least squares problem:

$$\hat{\mathbf{u}} = \arg\min_{\mathbf{u}} \|\mathbf{y} - \mathbf{HAu}\|_2^2 \tag{3.6}$$

Figure 3.1 illustrates the flowchart of the proposed stability-based SPFM algorithm.

3.3 Methods

The operation of the proposed stability-selection SPFM algorithm is illustrated in a dataset collected on a healthy subject in a 7T MR scanner (Siemens) using a 32-channel receive transmit coil. The subject was scanned under a Cleveland Clinic Institutional Review Board approved protocol (QED, Cleveland, OH). A volumetric MP2RAGE image was acquired for anatomical visualization. Two fMRI datasets were acquired with a simultaneous multislice EPI sequence (MB factor = 3, TE = 21 ms, field of view = 192×192 mm^2) at TR = 2800 ms ($1.2 \times 1.2 \times 1.5$ mm^3, flip angle = 55º) and 500 ms ($3 \times 3 \times 3$ mm^3, flip angle = 70º). For both TRs, the subject performed finger tapping events with the right index and thumb fingers every 45 s, where a single tap was performed in the first 6 minutes, or 10 taps quickly for the remaining 4 minutes. The onsets and durations of the paradigm are shown as gray vertical lines in Figure 3.2 (a) and Figure 3.3 (a).

Data preprocessing comprised an initial correction for motion using SLOMOCO2 (Beall & Lowe, 2014), detrending of 6th order Legendre polynomials and normalization to signal percentage change (SPC) with AFNI. Furthermore, a mask of white matter voxels was computed from the anatomical image with 3dSeg, which was then eroded 2 voxels to delimit voxels in deep white matter in the functional space. The preprocessed fMRI data were analyzed with three

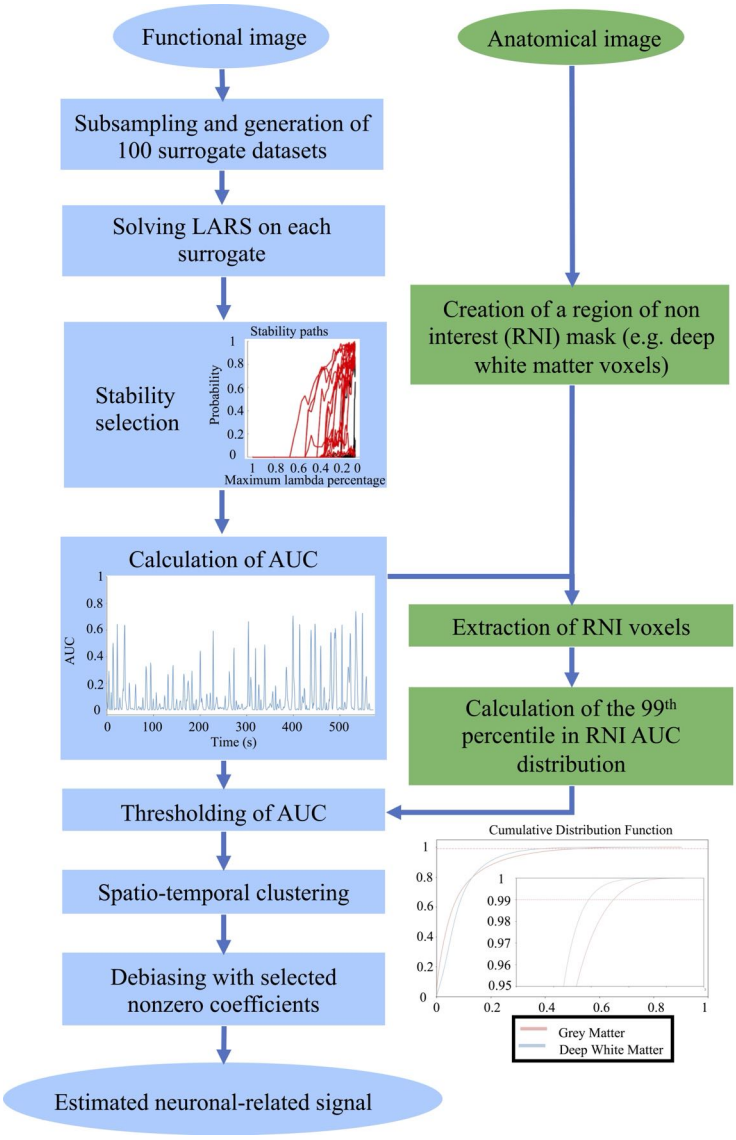

Figure 3.1: Flowchart of the stability-based SPFM algorithm.

different methods: 1) a traditional general linear model (GLM) analysis using the onsets and durations of the tapping events; 2) the original SPFM approach (3dPFM) using the LASSO for deconvolution and selection of the regularization parameter based on the Bayesian Information Criterion (BIC) (Gaudes et al., 2013); and 3) the novel stability-based SPFM with and without the integration operator in its formulation. Both SPFM approaches used the double-gamma

canonical HRF as a model for deconvolution (SPMG1 shape in 3dDeconvolve in AFNI). Previous to the final debiasing step, spatio-temporal clustering of a minimum of 5 contiguous voxels with activation (i.e., non-zero coefficient after thresholding) in a temporal window of ±1 TR was also performed to remove spurious, scattered activations.

3.4 Results and Discussion

Figure 3.2 and Figure 3.3 depict the activity maps estimated with all of the methods for different representative finger-tapping instants and the time courses of a voxel in the left primary motor cortex (marked with a white cross in the maps) for the high temporal and low spatial resolution dataset Figure 3.2 (a, b and c) and the low temporal and high spatial resolution dataset Figure 3.3 (a, b and c).

In the high temporal and low spatial resolution scenario (i.e., a high contrast-to-noise ratio regime), the activity maps in Figure 3.2 (c) illustrate that the original SPFM is able to detect finger tapping events with a high specificity. Implementing stability selection on the original SPFM algorithm increases the sensitivity while maintaining the specificity. However, as it can be seen in Figure 3.2 (a), the lack of an integration operator yields very variable estimates of the neuronal-related signal after debiasing with least squares (here scaled by 0.05 for visualization purposes) due to the large correlation of the debiasing model with contiguous non-zero coefficients at this fast temporal resolution. Conversely, the novel stability-based SPFM with the integration operator shows activity maps that are comparable to the ground truth despite the lower amplitude of the estimates. Yet, it can be observed that the signal model with the integrator overestimates the duration of the piecewise constant estimates for the short finger tapping events. Thus, in this scenario, the use of the stability selection and the innovation signal exhibits a similar performance to the original SPFM algorithm using LASSO and BIC since the high SNR and high temporal resolution (TR = 0.5 s) enables a precise and clear characterization of the dynamics of the BOLD signal, which facilitates the differentiation between noise and neuronal-related signal.

In an acquisition with a high spatial resolution and a low temporal resolution (i.e., a low contrast-to-noise ratio regime), Figure 3.3 (a), (b) and (c) demonstrate that the novel stability-based SPFM approach is able to detect more finger-tapping events and their associated brain activity than the original SPFM method. This advantage is clearly seen in the case of the single-tapping events, which exhibit a lower amplitude in the response than the long events with ten consecutive finger taps. The stability selection proves to be essential in correctly estimating finger tapping events, regardless of the use of the integration operator. The addition of the integration operator to the SPFM model produces activity maps that are closer to the ground truth of the GLM analysis (see Figure 3.2 c), even though the duration of the piecewise constant estimates is overestimated (see Figure 3.3 a). In this regime, the BIC criterion in the original SPFM is not

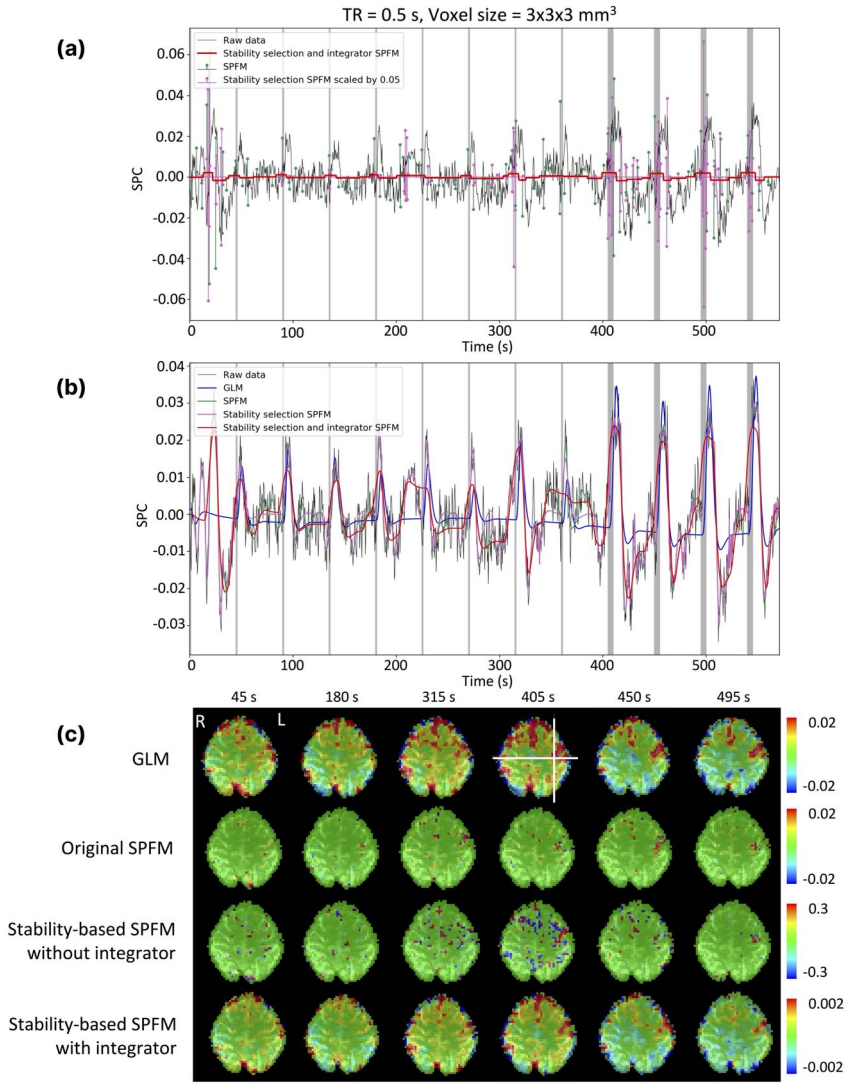

Figure 3.2: Comparison of the novel stability-based SPFM approach with the SPFM and the GLM methods for data with TR = 0.5 s and a voxel size of $3mm^3$ iso. (a) plots the time series of the voxels marked with a cross in (c) containing the raw data and the estimates of the different methods as shown in the legend. Onsets and duration of the finger-tapping are depicted as gray vertical lines. (b) shows the estimates of the different methods fitted with the canonical HRF. (c) shows the estimated maps of each of the methods for different finger-tapping events.

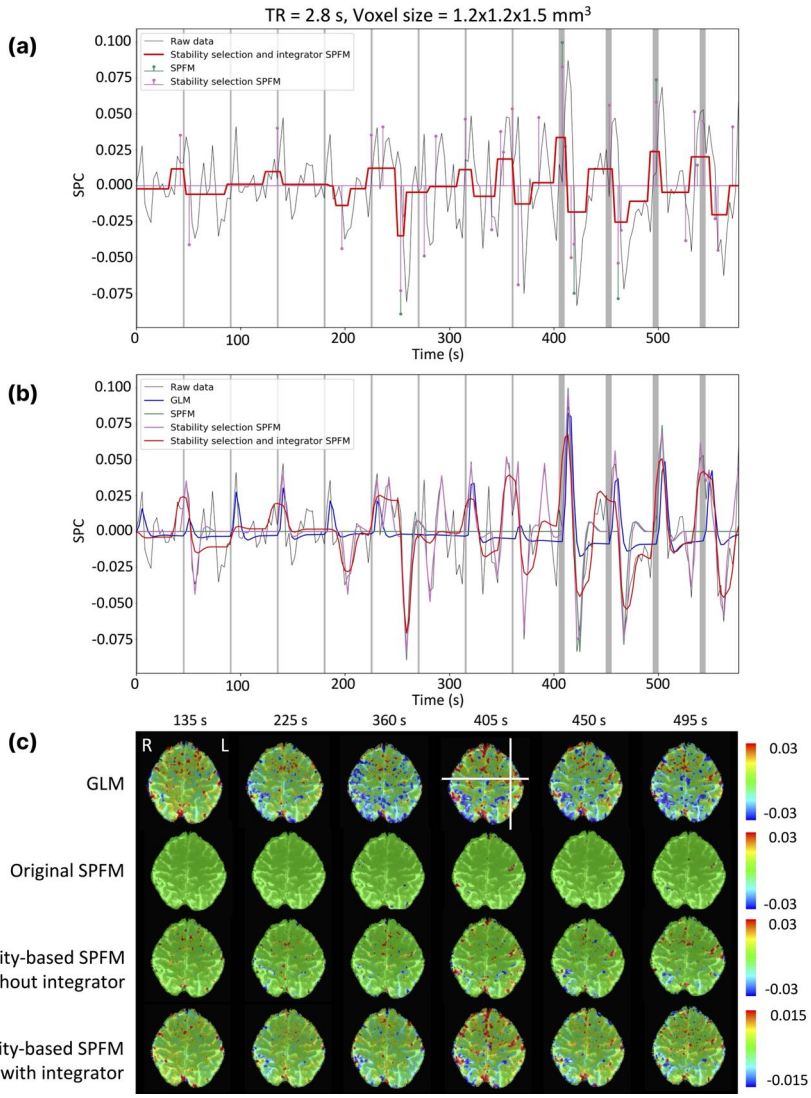

Figure 3.3: Comparison of the novel stability-based SPFM approach with the SPFM and the GLM methods for data with TR = 2.8 s and a voxel size of $1.2 \times 1.2 \times 1.5 mm^3$. (a) plots the time series of the voxels marked with a cross in (c) containing the raw data and the estimates of the different methods as shown in the legend. Onsets and duration of the finger-tapping are depicted as gray vertical lines. (b) shows the estimates of the different methods fitted with the canonical HRF. (c) shows the estimated maps of each of the methods for different finger-tapping events.

able to discern between neuronal-related events and noise, failing to detect the finger tapping events, probably as the shape of the BOLD response, which takes 4-6 s to reach its maximum amplitude, cannot be properly characterized by the model owing to the low temporal resolution (TR = 2.8 s). Hence, the stability selection procedure exhibits a robust performance at correctly estimating the neuronal-related events resulting from the finger tapping tasks, which showcases that the additions to the SPFM technique are promising, especially in low temporal resolution settings.

Chapter

4

Whole-Brain Multivariate Deconvolution for Multi-Echo Functional MRI

This chapter was published as Uruñuela, E., Gonzalez-Castillo, J., Zheng, C., Bandettini, P., & Caballero-Gaudes, C. (2022). Whole-brain multivariate hemodynamic deconvolution for multi-echo fMRI with stability selection. DOI: https://doi.org/10.1016/j.media.2023.103010.

This chapter proposes a novel hemodynamic deconvolution algorithm, multivariate sparse paradigm free mapping (Mv-SPFM), that operates at the whole brain level and adds spatial information via a mixed-norm regularization term over all voxels. Additionally, Mv-SPFM employs the stability selection procedure that removes the need to select regularization parameters and also lets us obtain an estimate of the true probability of having a neuronal-related BOLD event at each voxel and time-point based on the area under the curve (AUC) of the stability paths. Besides, its formulation is adapted for multi-echo fMRI acquisitions (MvME-SPFM), which allows us to better isolate fluctuations of BOLD origin on the basis of their linear dependence with the echo time (TE) and to assign physiologically interpretable units (i.e., changes in the apparent transverse relaxation ΔR_2^*) to the resulting deconvolved events. Remarkably, Mv-SPFM also achieves comparable performance when using a conventional single-echo formulation. The MV-SPFM algorithm outperforms previous hemodynamic deconvolution approaches, showing higher spatial and temporal agreement with the activation maps and BOLD signals obtained with a standard model-based linear regression approach. Furthermore, owing to the stability selection, the proposed algorithm provides more reliable estimates of neuronal-related activity for the study of the dynamics of brain activity when no information about the timings of the BOLD events is available. This algorithm is publicly available as part of the splora Python package.

4.1 Introduction

Conventionally, the analysis of functional MRI (fMRI) data relies on available information about the experimental paradigm to establish hypothesized models of brain activity. However, this information can be inaccurate, incomplete or unavailable in multiple scenarios such as resting-state, naturalistic paradigms or clinical conditions. In these cases, blind estimates of neuronal-related activity can be obtained with paradigm-free analysis methods such as hemodynamic deconvolution. Yet, current formulations of the hemodynamic deconvolution problem have three important limitations: 1) their efficacy strongly depends on the appropriate selection of regularization parameters, 2) being univariate, they do not take advantage of the information present across the brain, and 3) they do not provide any measure of statistical certainty associated with each detected event.

Despite the range of deconvolution methods that have been developed, few capitalize on the various properties of fMRI data, such as the advantages of multi-echo fMRI for denoising fMRI data (Bright & Murphy, 2013; Kundu et al., 2017), or the use of tissue-based or parcellation-based information to improve the accuracy of the estimates of neuronal activity. Recent exceptions include deconvolution algorithms that incorporate a multivariate formulation to perform spatio-temporal deconvolution (Bolton et al., 2019; Cherkaoui et al., 2021; Costantini et al., 2022; Uruñuela et al., 2021), thus addressing the second limitation stated in the previous paragraph. In addition, one deconvolution algorithm has been presented that exploits the mono-exponential decay model of the multi-echo fMRI signal: multi-echo sparse Paradigm Free Mapping (ME-SPFM) (Caballero-Gaudes et al., 2019). Furthermore, approaches have been developed to deal with the third limitation stated above and estimate the likelihood of having a neuronal event at each time-point and for each voxel by means of logistic regression (Bush & Cisler, 2013; Bush, Cisler, et al., 2015) or Gaussian mixture models (Pidnebesna et al., 2019). However, wouldn't it be desirable to have an algorithm that addresses all three limitations mentioned above? Specifically, it would be beneficial to obtain a measure of the probability of each voxel containing a neuronal event at each time-point using regularized estimators, while also leveraging the spatio-temporal information and physical properties of the fMRI signal for estimating the activity-inducing signal.

This chapter proposes a novel approach for the hemodynamic deconvolution of fMRI data that operates at the whole-brain level (i.e., multivariate formulation) to incorporate spatial information through a mixed-norm regularization term. Furthermore, this work proposes a stability selection procedure (Meinshausen & Bühlmann, 2010) that makes the estimation of the neuronal activity more robust to the selection of the regularization parameters, while providing the likelihood of having a neuronal-related event at each time-point and for each voxel. Using multi-echo fMRI data acquired from 10 healthy subjects (16 datasets) this chapter demonstrates that the proposed multivariate multi-echo Paradigm Free Mapping (MvME-SPFM) algorithm not only provides more robust estimates of the neuronal activity, but also yields a measure of the probability of each

voxel containing a neuronal event at each time-point. Moreover, MvME-SPFM returns quantitative estimates of ΔR_2^* in interpretable units (s^{-1}), which is relevant for functional analysis across different acquisition methods and field strengths. The chapter is structured as follows: the multivariate signal model and the multivariate multi-echo PFM algorithm are introduced in Section 4.2; Section 4.4 describes the data used and the analysis performed to evaluate this novel algorithm; the results of applying this new algorithm on experimental fMRI data are presented and compared to the previous PFM algorithms in Section 4.5; finally, the findings are discussed in Section 4.6.

4.2 Theory

4.2.1 Signal Model for Whole-Brain Deconvolution

Chapter 2 introduced that the BOLD fMRI signal model of a particular voxel can be written in matrix notation as:

$$\mathbf{y} = \mathbf{H}\Delta\mathbf{s} + \mathbf{e} \tag{4.1}$$

where, assuming zero-boundary conditions, the variables $\mathbf{y}$, $\Delta\mathbf{s}$, $\mathbf{e} \in \mathbb{R}^N$ are the voxel's time-series, the activity-inducing signal changes and the noise term, respectively, and $\mathbf{H} \in \mathbb{R}^{N\times N}$ is the Toeplitz convolution matrix defined by the HRF (Gaudes et al., 2013; Gitelman et al., 2003). Hereinafter in the chapter, the noise term is assumed to be an additive white Gaussian noise (i.e., uncorrelated). Furthermore, the $\Delta\mathbf{s}$ can be defined in a finer temporal resolution (i.e., increasing its dimension by a factor $\alpha > 1$ such that $\Delta\mathbf{s} \in \mathbb{R}^{\alpha N}$) so that it can describe activity-inducing signal changes that are asynchronous to the TR of the data. Though in such scenario, the convolution matrix becomes non-Toeplitz (i.e. $\mathbf{H} \in \mathbb{R}^{N\times \alpha N}$) (Ciuciu et al., 2003).

The previous signal model is applicable for single-echo and multi-echo fMRI data. In the particular case of multi-echo fMRI acquired with a gradient-echo sequence, the voxel's time-series in terms of the signal percentage change has a linear relationship with the echo time (TE) as $y(\mathrm{TE}_k, n) \approx \Delta\rho(n) - \mathrm{TE}_k \Delta R_2^*(n)$, where $\Delta R_2^*(n)$ denotes the BOLD-like signal changes and $\Delta\rho(n)$ corresponds to changes in the net magnetization, for instance due to head motion (Kundu et al., 2017). The signal changes associated to fluctuations in the net magnetization can be effectively reduced in preprocessing, for example using multi-echo independent component analysis (Caballero-Gaudes et al., 2019; Kundu et al., 2012), and are neglected hereinafter. Hence, considering that neuronal-related signal changes $\Delta\mathbf{s}$ produce a change in ΔR_2^*, the signal model in Eq.(4.1) can be adapted to contain the signal acquired at all K echo-times (TE) via concatenation:

$$\begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_K \end{bmatrix} = - \begin{bmatrix} \mathrm{TE}_1\mathbf{H} \\ \vdots \\ \mathrm{TE}_K\mathbf{H} \end{bmatrix} \Delta\mathbf{s}, \tag{4.2}$$

which can be simplified into $\bar{\mathbf{y}} = -\overline{\mathbf{H}}\Delta\mathbf{s}$, where $\bar{\mathbf{y}} \in \mathbb{R}^{KN}$, $\Delta\mathbf{s} \in \mathbb{R}^N$ and $\overline{\mathbf{H}} \in \mathbb{R}^{KN\times N}$.

Assuming that the shape of the hemodynamic response can be similarly modeled across all brain voxels, the previous voxel-wise (i.e., univariate) model in Eq.(4.2) can be extended straightforwardly to a multivariate formulation that considers all the voxels V of the brain:

$$\begin{bmatrix} \mathbf{y}_{1,1} & \cdots & \mathbf{y}_{1,V} \\ \vdots & \ddots & \vdots \\ \mathbf{y}_{K,1} & \cdots & \mathbf{y}_{K,V} \end{bmatrix} = - \begin{bmatrix} \mathrm{TE}_1\mathbf{H} \\ \vdots \\ \mathrm{TE}_K\mathbf{H} \end{bmatrix} \begin{bmatrix} \Delta\mathbf{s}_1 & \cdots & \Delta\mathbf{s}_V \end{bmatrix}, \tag{4.3}$$

which can be simplified into $\bar{\mathbf{Y}} = -\bar{\mathbf{H}}\Delta\mathbf{S}$, where $\bar{\mathbf{Y}} \in \mathbb{R}^{KN\times V}$, $\bar{\mathbf{H}} \in \mathbb{R}^{KN\times N}$ and $\Delta\mathbf{S} \in \mathbb{R}^{N\times V}$.

4.3 Solving the Optimization Problem for Whole-Brain Deconvolution

As described in Chapter 2, an estimate of the activity that induces the BOLD response $\hat{\mathbf{s}}$ can be obtained by solving an ordinary least-squares problem such as:

$$\Delta\hat{\mathbf{s}} = \arg\min_{\mathbf{s}} \frac{1}{2}\|\bar{\mathbf{y}} - \bar{\mathbf{H}}\Delta\mathbf{s}\|_2^2 + \lambda\|\Delta\mathbf{s}\|_1, \tag{4.4}$$

where λ is the regularization parameter that regulates the level of sparsity of the estimates given the ℓ_1-norm, which is defined as $\|\Delta\mathbf{s}\|_1 = \Sigma_{n=1}^{N}|\Delta\mathbf{s}_n|$.

The inverse problem in Eq.(4.4) can be directly adapted to be solved at the whole-brain using the multivariate formulation in Eq.(4.3). More interestingly though, solving the inverse problem at the whole-brain level opens up many possibilities in the form of additional regularization terms to take advantage of the spatial information for an informed estimation of the activity-inducing signal $\Delta\hat{\mathbf{S}}$. For instance, mixed-norms in the form of $\ell_{p,q}$ can be employed to separate coefficients into groups that are blind to each other, while the coefficients within a group are treated together (Kowalski, 2009). Hence, regularization terms based on mixed-norms can promote spatio-temporal structures that are observed in fMRI signals.

Here, an $\ell_{2,1} + \ell_1$ mixed-norm regularization term (Gramfort et al., 2011) is added to the multivariate convex problem to promote the co-activation of the activity-inducing signal $\Delta\hat{\mathbf{S}}$ considering the coefficients of the voxels of the brain (V) at time n as one group:

$$\Delta\hat{\mathbf{S}} = \arg\min_{\mathbf{S}} \frac{1}{2}\|\bar{\mathbf{Y}} - \bar{\mathbf{H}}\Delta\mathbf{S}\|_2^2 + \lambda\rho\|\Delta\mathbf{S}\|_1 + \lambda(1-\rho)\|\Delta\mathbf{S}\|_{2,1}, \tag{4.5}$$

where $\ell_{2,1}$-norm is defined as $\|\Delta\mathbf{S}\|_{2,1} = \Sigma_{n=1}^{N}\sqrt{\Sigma_{v=1}^{V}\Delta\mathbf{S}_{n,v}^2}$, and $0 < \rho < 1$ is a parameter that controls the tradeoff between the sparsity introduced by the ℓ_1-norm and the grouping of voxels promoted by the $\ell_{2,1}$-norm so that the estimation of one voxel coefficient at time n is influenced by the estimates of the rest of the brain voxels at the same time. Note that when $\rho = 1$ Eq. (4.5) is the whole-brain equivalent of Eq. (4.4). Additionally, it is worth noting that

mixed-norm regularization, such as group LASSO, has been previously used in MEG/EEG studies (Gramfort et al., 2011). However, in those cases, the spatial and temporal dimensions were swapped compared to Eq. (4.5).On the other hand, the regularization parameter λ can be adapted voxel-wise in order to account for differences in the signal-to-noise ratio across voxels. Consequently, the multivariate deconvolution problem can be written as:

$$\Delta\hat{\mathbf{S}} = \arg\min_{\mathbf{S}} \frac{1}{2}\|\bar{\mathbf{Y}} - \bar{\mathbf{H}}\Delta\mathbf{S}\|_2^2 + \rho\|\mathbf{D}\Delta\mathbf{S}\|_1 + (1-\rho)\|\mathbf{D}\Delta\mathbf{S}\|_{2,1}, \tag{4.6}$$

where $\mathbf{D} = \mathrm{diag}(\lambda_1, \ldots, \lambda_V) \in \mathbb{R}^{V\times V}$ is a diagonal matrix with the voxel-specific values of λ. In practice, a criterion must be used to select the voxel-specific λs. Instead, the use of stability selection to avoid this critical choice is proposed (see Section 4.4.2).

Therefore, given the convex nature of the inverse problem in Eq. (4.6), estimates of $\Delta\hat{\mathbf{S}}$ can be calculated using the fast iterative shrinkage-thresholding algorithm (FISTA) (Beck & Teboulle, 2009) with the following proximity operator for $\ell_1 + \ell_{2,1}$:

$$S_{n,v} = \frac{Z_{n,v}}{|Z_{n,v}|}\left(|Z_{n,v}| - \lambda_v\rho\right)^+ \left(1 - \frac{\lambda_v(1-\rho)}{\sqrt{\sum_v \left(|Z_{n,v}| - \lambda_v\rho\right)^{+2}}}\right)^+, \tag{4.7}$$

where $\Delta\mathbf{S} = \mathrm{prox}_{\lambda(\rho\|\cdot\|_1 + (1-\rho)\|\cdot\|_{2,1})}(\mathbf{Z}) \in \mathbb{R}^{N\times V}$, $(x)^+ = \max(x, 0)$ for $x \in \mathbb{R}$, and $\frac{0}{0} = 0$ by convention.

4.4 Methods

4.4.1 fMRI Data Acquisition and Preprocessing

The evaluation of the proposed MvME-SPFM was performed on ME-fMRI data acquired in 10 subjects using a multi-task rapid event-related paradigm. Six subjects performed two functional runs, the other 4 subjects only performed 1 run due to scanning time constraints (i.e., a total of 16 datasets). All participants gave informed consent in compliance with the NIH Combined Neuroscience International Review Board-approved protocol 93-M-1070 in Bethesda, MD. A thorough description of the MRI acquisition protocols and experimental tasks in the experimental design can be found in (Gonzalez-Castillo et al., 2016), only those details that are relevant to this analysis are given here.

MRI data was acquired on a General Electric 3T 750 MRI scanner with a 32-channel receive-only head coil (General Electric, Waukesha, WI). Functional scans were acquired with a ME gradient-recalled echoplanar imaging (GRE-EPI) sequence (flip angle = 70° for 9 subjects, flip angle = 60° for 1 subject, TEs = 16.3/32.2/48.1 ms, TR = 2 s, 30 axial slices, slice thickness = 4 mm, in-plane resolution = 3×3 mm^2, FOV 192 mm, acceleration factor 2, number of acquisitions = 220). Functional data was acquired with ascending sequential slice acquisitions, except in one subject where the acquisitions were interleaved.

In addition, high resolution T1-weighted MPRAGE and proton density images were acquired per subject for anatomical alignment and visualization purposes (176 axial slices, voxel size $= 1 \times 1 \times 1$ mm^3, image matrix $= 256 \times 256$).

Each run of data acquisition consisted of 6 trials with 5 different tasks each: biological motion observation (BMOT), finger tapping (FTAP), passive viewing of houses (HOUS), listening to music (MUSI), and sentence reading (READ). The reader is referred to that paper for details on the preprocessing steps, and comparison with alternative single-echo models for deconvolution. This data had previously been employed, preprocessed and ME-ICA denoised for the evaluation of the ME-SPFM algorithm in (Caballero-Gaudes et al., 2019).

4.4.2 Stability Selection of the Regularization Parameter λ

The choice of the regularization parameter λ is crucial to obtain accurate estimates of $\Delta\hat{\mathbf{S}}$. Although the value of λ of each voxel could be fixed ad-hoc, previous work has opted for the use of model selection criteria, such as the Bayesian Information Criterion (BIC), on the regularization path (Caballero-Gaudes et al., 2019), computed by means of the least angle regression (LARS) algorithm (Efron et al., 2004). Even though the use of BIC performed well for ME-SPFM (Caballero-Gaudes et al., 2019) and its single-echo counterpart (SPFM) (Gaudes et al., 2013), due to its high specificity, it can be problematic for certain voxels where the BIC curve might present multiple local minima or even fail to present a clear minima for the evaluated range of λ.

This chapter proposes a more robust procedure to address this shortcoming for selecting λ with the usage of the stability selection method (Meinshausen & Bühlmann, 2010) described in Chapter 3. Moreover, the stability selection procedure presented here yields the probability to have a non-zero coefficient in the activity-inducing signal at each time-point. Specifically, our implementation of the stability selection procedure generates $T = 30$ surrogates of each voxel timecourse by randomly subsampling 60% of the time-points (a more computationally expensive version was also tested with $T = 100$ surrogates that yielded very similar results). The convolution matrix $\mathbf{H}$ is subsampled accordingly. The surrogate data is then analyzed with the inverse problem in Eq. (4.5) for a voxel-specific range of λ values using the fast iterative shrinkage thresholding algorithm (FISTA). Here, 30 values of λ are spaced logarithmically between 5% and 95% of the voxel-specific maximum λ possible to more accurately sample the lower range. Then, the ratio (probability) of surrogates where the estimated coefficient at each time-point is non-zero is calculated for each time-point and value of λ. As illustrated in Figure 4.1, these probabilities build the so-called stability paths, which resemble the well-known regularization paths of conventional regularized estimators (e.g., LASSO, Ridge Regression) that plot the amplitude of the coefficients for each λ.

Unlike the original stability selection procedure, which sets a given probability threshold to select the final set of non-zero coefficients (Meinshausen & Bühlmann, 2010), the area under the curve (AUC) of the stability path of each time-point is calculated as an index of confidence of having a non-zero coefficient across

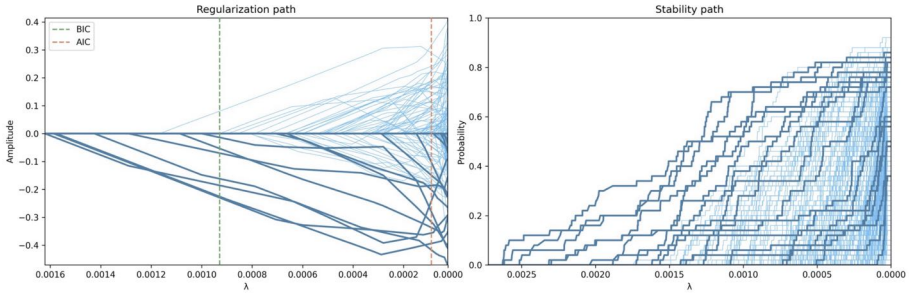

Figure 4.1: Example of the regularization path and the stability path for a voxel timeseries with $\rho = 1$. On the left, the regularization path shows the amplitude of each coefficient estimate $\Delta\hat{\mathbf{S}}$ (one per TR). At first, all the coefficients are zero and successively they become non-zero as λ decreases towards zero, which corresponds to the orthogonal least squares solution (i.e., no regularization). On the right, the corresponding stability path plots the probablity that each coefficient estimate (i.e., each TR) is non-zero for each value of λ based on the stability selection procedure. Note that both paths can have a different maximum value of λ given the subsampling step in the stability selection. Lines in the stability path correspond to different TRs of a single voxel. The darker lines denote the coefficient estimates corresponding to the TRs during the task-related events.

the evaluated range of λ. As a result, the AUC timecourse provides a measure of the probability of having neuronal-related activity at each time-point and voxel. Next, the AUC time-series are thresholded according to the histogram of AUC values in a region of non-interest (hereinafter, denoted as the null AUC histogram) to yield a sparse representation of the signal. Alternatively, a null distribution of AUC values could be generated from surrogate data (Liégeois et al., 2021). Accordingly, when employing stability selection, the individual voxels' estimates might not be equivalent to the voxels' estimates in any single one of the whole-brain models that can be formulated with a given value of λ in Eq.(4.5) or $\mathbf{D}$ in Eq. (4.6)), but are rather obtained by computing area-under-the-curve (AUC) values for neuronal-related events.

Finally, a fitting step is applied to each voxel by defining a reduced convolution model with the selected non-zero coefficients and fitting it by means of a conventional orthogonal least squares estimator. This step reduces the bias towards zero imposed by the sparsity-promoting regularization terms, and thus obtains more realistic estimates of the neuronal-related signal (here, in terms of ΔR_2^*) (Caballero-Gaudes et al., 2019).

4.4.3 Balancing the Spatial Regularization

The $\ell_{2,1}$-norm regularization term in Eq. (4.5) promotes structured spatio-temporal sparsity in the sense that the estimates of all brain voxels at a given time-point are treated as a group and this term forms a constraint on the number of groups with at least one non-zero estimate to model the data. Assuming that $\rho = 0$, either the value of all the voxel estimates at one time point can be non-zero or all of them are nulled. Hence, this regularization term considers spatial information from all brain voxels for the deconvolution since the value of a given voxel coefficient also depends on the rest of the voxels.

To illustrate the effect of the corresponding regularization parameter ρ, this chapter solves the multivariate regularization problem in Eq. (4.5) using stability selection for $\rho = 1$, $\rho = 0.5$ and $\rho = 0.$; i.e., applying the sparsity-promoting ℓ_1-norm only, equally weighting the sparsity and spatial regularizations, and employing the $\ell_{2,1}$-norm spatial regularization only, respectively.

4.4.4 Comparison with Conventional Timing-Based GLM Analyses

To evaluate how the multivariate formulation combined with stability selection improves the accuracy of the estimates of $\Delta\hat{\mathbf{S}}$ compared with its univariate counterpart ME-SPFM using the BIC for voxel-wise selection of λ (Caballero-Gaudes et al., 2019), the spatial sensitivity, specificity and overlap (using a Dice coefficient metric) of the MvME-SPFM activation maps were calculated using the trial-level GLM-based activation maps ($p \leq 0.05$) as the ground truth. These analyses were computed with the independently modulated (IM) option of the 3dDeconvolve program in AFNI that implements an orthogonal least squares estimation for each trial, thus also assuming uncorrelated noise as in our model. Single trial GLM maps (GLM-IM) were obtained from the optimally combined and ME-ICA denoised data, and only negative ΔR_2^* (i.e., $\Delta\hat{\mathbf{S}} < 0$ that generate a positive BOLD response) were considered for the computation of the Dice coefficients.

For the MvME-SPFM, the following two strategies for thresholding the AUC timeseries were considered in order to define the corresponding activation maps:

- Static thresholding (ST): The estimates of $\Delta\hat{\mathbf{S}}$ obtained with the novel MvME-SPFM technique that utilizes stability selection, where the AUC threshold was chosen as the 95th percentile of the histogram of AUC in deep white matter voxels (i.e., a fixed, static threshold), which were labeled after tissue segmentation of the T1-weighted anatomical MRI using 3dSeg in AFNI, and eroding 4 voxels of the resulting white matter tissue mask at anatomical resolution.

- Time-dependent thresholding (TD): The estimates of $\Delta\hat{\mathbf{S}}$ obtained with the novel MvME-SPFM technique with stability selection, where the AUC threshold varies temporally according to the 95th percentile of the null histogram of AUC at each time-point. This implementation was based on the hypothesis that a time-dependent threshold would be able to better

control for widespread spurious deconvolved changes in $\Delta\hat{\mathbf{S}}$, for instance due to head motion or deep breaths.

Note the ST and TD thresholding strategies can be applied in the analysis of single-echo and multi-echo fMRI data.

4.4.5 Comparison with Another State-of-the-Art Multivariate Deconvolution Method

In order to evaluate how the MvME-SPFM algorithm compares with another state-of-the-art multivariate deconvolution method, Hemolearn (Cherkaoui et al., 2021) was chosen, a recently proposed multivariate semi-blind deconvolution method that can estimate the HRF and the neuronal signal in a paradigm-free setting and that has been shown to faithfully capture intrinsic functional connectivity networks at the subject level. This technique introduces a low-rank constraint to learn both K temporal atoms (with $K \ll P$) and their corresponding spatial maps, which encode various functional networks, each of them with their specific neural activation profile. Hence, HemoLearn models the brain activity as a linear combination between the neural activations $Z = (z_k)_{k=1}^{K} \in \mathbb{R}^{K \times T}$ and the spatial patterns $U = (u_k)_{k=1}^{K} \in \mathbb{R}^{K \times P}$, where T is the total number of TRs in the data and P is the number of voxels. The BOLD fMRI signal model is then given by the convolution between an estimated regional HRF and the activation as:

$$\mathbf{Y} = \left(\sum_{m=1}^{M} \mathbf{H_m} \right) * \left(\sum_{k=1}^{K} \mathbf{u_k}^T \mathbf{z_k} \right) + \mathbf{E} \tag{4.8}$$

Thus, the minimization problem proposed by HemoLearn can be described as:

$$\arg\min_{\mathbf{U},\mathbf{Z}} = \frac{1}{2} \| \mathbf{Y} - \left(\sum_{m=1}^{M} \mathbf{H_m} \right) * \left(\sum_{k=1}^{K} \mathbf{u_k}^T \mathbf{z_k} \right) \|_F^2 + \lambda \sum_{k=1}^{K} \| \nabla \mathbf{z_k} \|_1 \tag{4.9}$$

The practical implementation of Hemolearn –which is available at https://hemolearn.github.io/– requires the specification of the number of components to be estimated, which has no standard procedure to be determined. Here, the analysis was performed with a range of components from 5 to 40. With the aim of making the comparison as fair as possible, the component that maximized the correlation between the estimated activity map and the session-level GLM-based activation map was selected for each condition in the task. The selected component time-series were compared with the corresponding $\Delta\hat{\mathbf{S}}$ obtained with MvME-SPFM using stability selection and the time-dependent thresholding (TD), and the selected component activity maps with the corresponding trial-level MvME-SPFM activity maps obtained with the same thresholding strategy.

4.5 Results

The output of deconvolution algorithms such as ME-SPFM and the proposed MvME-SPFM is a 4D dataset that matches the dimensions (both spatial and temporal) of the input data, i.e., it is a movie of the estimated ΔR_2^* maps. In addition, the use of stability selection generates the area under the curve (AUC) 4D output dataset, which indicates the probability of having a neuronal-related event at each time-point for every voxel in the brain.

Figure 4.2 depicts the area under the curve (AUC) time-series and maps obtained with stability selection for $\rho = 0.5$ in representative voxels of each task in the paradigm (indicated with a cross in the maps), where the AUC maps correspond to single time-points signaled by the blue arrows. The AUC time-series of the ST and TD thresholding approaches are shown on top of the original AUC time-series. The AUC maps depict spatial patterns of ΔR_2^* where regions that are typically involved in the tasks show higher probabilities of having neuronal-related activity compared with other brain regions.

Figure 4.3 displays the comparison of the ΔR_2^* maps obtained by solving the inverse problem in Eq. (4.5) with a fixed selection of λ (1^{st} row) and with the use of stability selection (2^{nd}, 3^{rd} and 4^{th} rows) for $\rho = \{0, 0.5, 1\}$. The ΔR_2^* maps obtained with a fixed selection of λ equal to the noise estimate of the first echo volume (1^{st} row) are very sensitive to the selection of ρ. Similar observations were obtained with other values of λ. With a selection of $\rho = 1$, only the ℓ_1-norm regularization term is applied, which produces ΔR_2^* maps with few non-zero coefficients. With $\rho = 0$, only the $\ell_{2,1}$-norm spatial regularization is applied, which yields a ΔR_2^* map that covers the entire brain and does not exhibit a spatial pattern in concordance with the task. However, a selection of $\rho = 0.5$ yields a ΔR_2^* map that is more similar to the activity maps often observed when participants are asked to look at the image of a house, depicting negative ΔR_2^* in bilateral fusiform regions. In contrast, the use of stability selection yields AUC maps (row 2) and the corresponding ΔR_2^* maps after each thresholding strategy (rows 3-4) reveal activation patterns concordant with those often seen for viewing houses regardless of the selection of ρ. In other words, the ΔR_2^* maps obtained with stability selection are less sensitive to the selection of ρ while obviating the need to choose λ. In fact, the spatial correlations between the AUC maps for each pair of ρ's were nearly equal to 1 for all time points (average correlations are 0.97 between $\rho = \{0, 0.5\}$, 0.98 between $\rho = \{0, 1\}$, and 0.97 between $\rho = \{0.5, 1\}$). In addition, it can be seen that using a TD threshold yields BOLD signal changes that are more confined to the expected areas in bilateral fusiform cortices than the ST threshold. Due to the high similarity of the AUC maps for any value of ρ, only the results for $\rho = 0.5$ are discussed hereinafter.

Figure 4.4 provides an in-depth view of how the time-dependent thresholding operates when motion- and respiration-related artifacts are present in the data. The grayplot (Power, 2017) in Figure 4.4A clearly shows bands spanning throughout the entire brain that illustrate significant changes in the amplitude of the signal. The source of these signal changes can be attributed to head

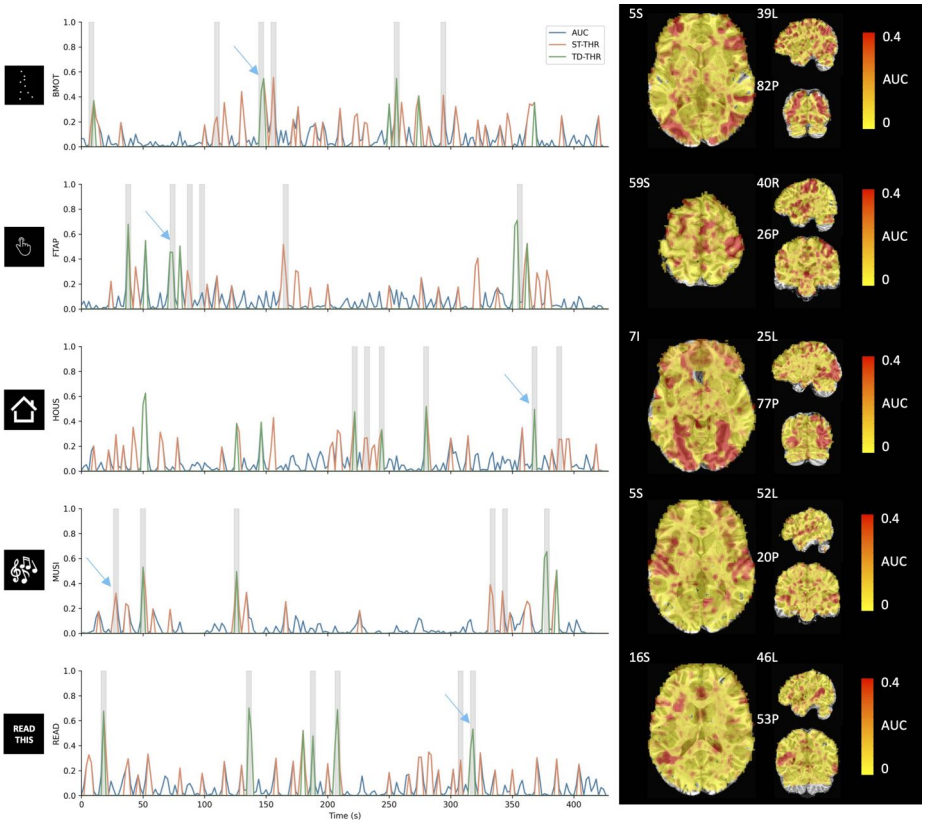

Figure 4.2: Left: Original (blue), ST thresholded (orange) and TD thresholded (green) AUC time-series for a representative voxel for each task in the paradigm ($\rho = 0.5$). Note that the three time-series are overlaid; i.e., the static and time-dependent time-courses are thresholded versions of the original AUC. Gray blocks depict the onset and duration of each trial. Right: AUC maps at the time-points signaled by the blue arrows.

motion events (see Euclidean norm in Figure 4.4C) and deep breaths (see arrows for respiration volume signal (Chang et al., 2009) in Figure 4.4D). The respiration-related events cause a drop in the global signal (see Figure 4.4B) seconds after the peak in the respiration volume signal. Interestingly, our results show a decrease in the equivalent ST percentile that corresponds to the 95th TD threshold (Figure 4.4E) at the instances of these large respiratory-related events. This decrease can also be observed in the corresponding AUC value of the TD thresholding strategy as shown in Figure 4.4F. The distributions of AUC values at the time-points with respiratory- and motion-related artifacts have a shorter tail than the distribution of the AUC values at the time-points where subjects performed the task. Hence, in these events the TD thresholding strategy is able

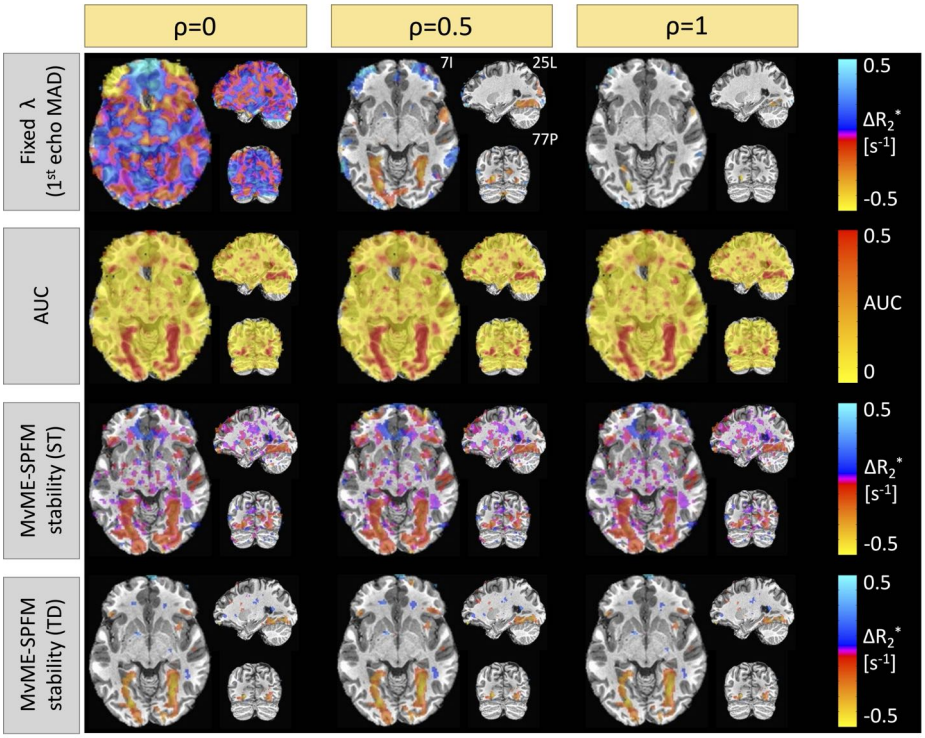

Figure 4.3: Comparison of the ΔR_2^* maps obtained with a fixed selection of λ (row 1) and the use of stability selection (rows 2-4: AUC, stability selection with static thresholding (ST), and stability selection with time-dependent thresholding (TD)) for $\rho = 0$ (column 1), $\rho = 0.5$ (column 2), and $\rho = 1$ (column 3). These maps correspond to a single-trial event of the house-viewing task (HOUS).

to adjust the threshold so that the final estimates of ΔR_2^* specifically capture task-activated voxels while excluding voxels that are affected by artifacts. The higher specificity of the TD thresholding strategy can be clearly seen in the ROC values shown in Figure 4.4H-L. The use of stability selection with the TD threshold yields more specific estimates of ΔR_2^* than with ST thresholding or the original ME-SPFM method, while the sensitivity is slightly reduced. On the other hand, the use of stability selection with a ST threshold improves the sensitivity of the ΔR_2^* estimates compared to the original ME-SPFM technique while preserving its specificity.

Figure 4.5 depicts the time-series of the estimated ΔR_2^* and denoised BOLD, i.e., ΔR_2^* convolved with the HRF, for a representative voxel of each task for the subject depicted in Figure 4.4 and Figure 4.7 and compared to a reference voxel in the lateral ventricles. The location of the voxels is shown in the corresponding maps in Figure 4.7. The ST thresholding approach detects ΔR_2^* events of the

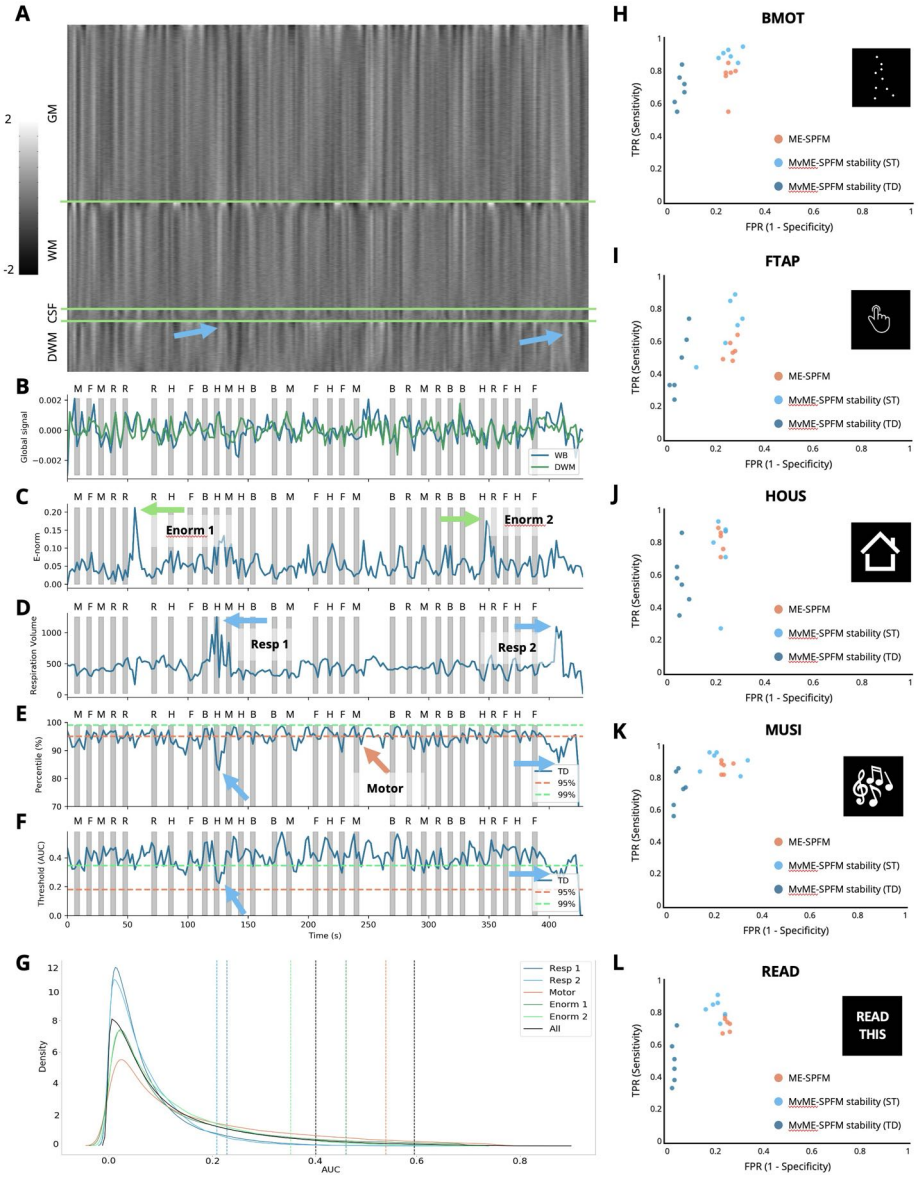

Figure 4.4: Representative subject with motion and respiration artifacts. A: Grayplot of the second echo volume. The grayplot is divided into 4 sections: gray matter (GM), white matter (WM), cerebrospinal fluid (CSF), and deep white matter (DWM). B: Time-series of the global signal calculated in the whole brain (WB, blue) and the deep white matter (DWM, green).

Figure 4.4: C: Euclidean norm (e-norm) of the temporal derivative of the realignment parameters. D: respiration volume signal. E: AUC percentile corresponding to the time-dependent threshold (lines at 95th and 99th percentiles are shown for reference).F: AUC values corresponding to the time-dependent threshold are shown in blue. The horizontal dashed lines indicate the 95th (orange) and 99th (green) percentiles corresponding to ST thresholding. Gray bars in B-F indicate the onset and duration of each trial in the paradigm, with their respective initials on top. Blue arrows point out two respiration-related events, green arrows point out two motion-related events, and the orange arrow points out a finger-tapping event. G: Probability density functions (estimated by kernel density estimate) of the AUC values corresponding to the instances of the two respiratory-related events (blue lines), a representative time-point of one finger-tapping trial (orange line), the two largest peaks in the e-norm trace (green lines), and the overall AUC distribution (black). The corresponding coloured vertical dashed lines indicate the AUC value for the 95th percentile of the TD thresholding approach, along with the 95th and 99th AUC values of ST thresholding. H-L: Receiver operating characteristic (ROC) values for the original ME-SPFM (orange), and proposed MvME-SPFM technique with the use of stability selection with the ST (light blue) and TD (dark blue) thresholding approaches for this dataset. The ROC plots depict the sensitivity and specificity of the methods at correctly estimating the activity maps that correspond to the 6 trials of each task in the paradigm.

activity-inducing signal that correctly match the timings of the stimuli (i.e., high temporal sensitivity), but also shows events that occur in the resting state and do not coincide with any activity-evoking trial. Based on comparison with the events detected in the time series extracted from the lateral ventricles, it can be conjectured that some of these events might be due to artifactual and physiological fluctuations that remain in the signal after preprocessing. On the other hand, ΔR_2^* values estimated with the TD thresholding approach match the timings of the stimuli almost perfectly with few missed trials (high temporal specificity). This is supported by the few ΔR_2^* events obtained for the reference voxel in the ventricles. Likewise, the denoised BOLD time-series obtained with the TD thresholding approach clearly describes signal changes associated with the trials, whereas the denoised BOLD time-series estimated with the ST thresholding strategy fits the original data very closely, which could be interpreted as a signature of overfitting.

Figure 4.7 illustrates the activation maps of representative single-trial events of each task for the same subject depicted in Figure 4.4. The activation maps of the proposed MvME-SPFM formulation were compared using the two thresholding approaches with the activation maps obtained with a single-trial GLM and the previous ME-SPFM approach. While all PFM methods exhibit activation maps that highly resemble those obtained with the single-trial GLM analysis, differences between the methods can be observed. For instance, although the use

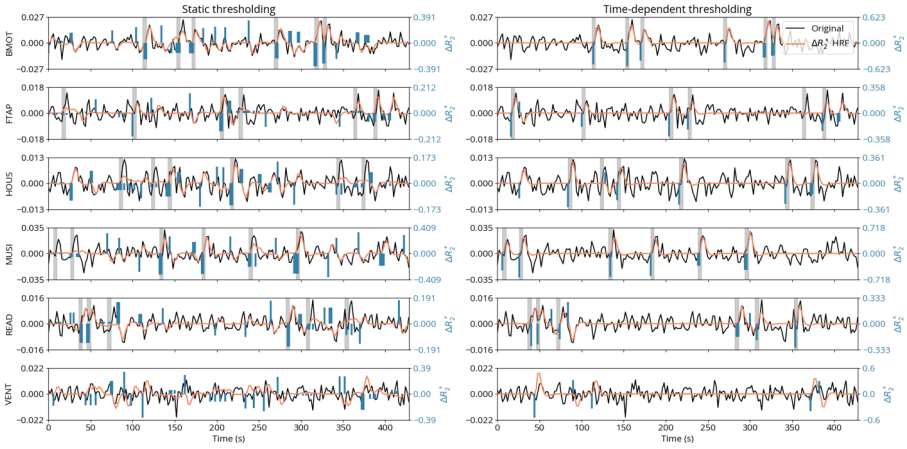

Figure 4.5: Comparison of the estimated ΔR_2^* (blue) and denoised BOLD (orange), i.e., ΔR_2^* convolved with the HRF, time-series when employing the ST (left) and TD (right) thresholding approches, for representative voxels of each task (rows) as well as one voxel from the lateral ventricle for reference. The estimates shown here were obtained with $\rho = 0.5$. The preprocessed time series is shown in black. The gray bars indicate the onset and duration of each trial for each task of the experimental paradigm.

of stability selection with a ST thresholding approach yields maps with clusters of activation of comparable size and location to those found with ME-SPFM, in certain noisy trials (e.g., see HOUS-Trial 1), the ST-thresholding MvME-SPFM maps can yield reduced spatial specificity, probably related to spurious, scattered changes in R_2^*. Across all tasks, the maps obtained with TD thresholding exhibit a notably larger resemblance to the single-trial GLM, showing higher spatial specificity and lower sensitivity compared to the other two PFM methods.

To further evaluate the proposed MvME-SPFM, the motor task of a single subject (100206) extracted from the Human Connectome Project (HCP) dataset (Van Essen et al., 2013) was analyzed. The data was 3 min 24 s long (after removing the first 10 seconds for steady magnetization) with a TR of 0.72 s, a multi-band factor of 8 and a spatial resolution of $2 \times 2 \times 2$ mm^3. The images were already preprocessed using a standard HCP pipeline including realignment, coregistration, spatial normalization and smoothing. The task was composed of 5 blocks of 12 s each, preceded by a 3 s cue indicating the task to be performed by the participant. The conditions in each 5-second block were: left hand finger tapping, right hand finger tapping, left foot movement, right foot movement, and tongue movement. The process was repeated once more, resulting in a total of 10 blocks. We adjusted the signal model in Eq. (4.3) to be used with single-echo data as described in (Gaudes et al., 2013; Uruñuela et al., 2023)

The accurate performance of the proposed MvME-SPFM method observed in

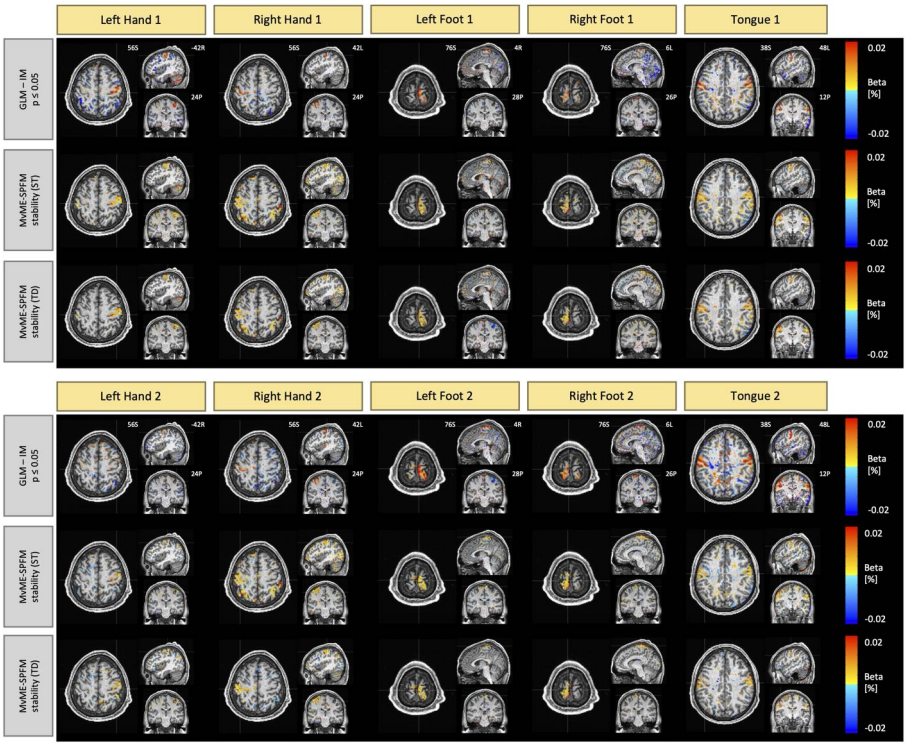

Figure 4.6: Single-trial activity maps estimated for the single-echo motor task data from the Human Connectome Project.Row 1 depicts the activation maps obtained with a single-trial GLM ($p \leq 0.05$), row 2 depicts the maps detected with the novel Mv-SPFM technique with stability selection, $\rho = 0.5$ and a static threshold (ST), and row 3 illustrates the results using a time-dependent threshold (TD). The activity maps corresponding to the first trial are shown on top, and the activity maps corresponding to the second trial are shown on the bottom.

the multi-echo fMRI data is consistent with the results found in the single-echo data from the Human Connectome Project. Figure 4.6 depicts the single-trial activity maps obtained with a GLM and the proposed approach with stability selection and both ST and TD thresholding. In general, the activity maps obtained with the proposed method are highly comparable to the single-trial GLM activation maps. Importantly, the proposed method showed a larger sensitivity to detect the activity evoked by the motor task for certain trials than the timing-aware GLM analysis. For instance, the Mv-SPFM activation maps of the second trial for the left hand finger tapping condition shows more BOLD activity in motor regions of the right precental gyrus, which is barely seen in the corresponding single-trial GLM map.

As illustrated in Figure 4.8, the Dice coefficient between the estimated

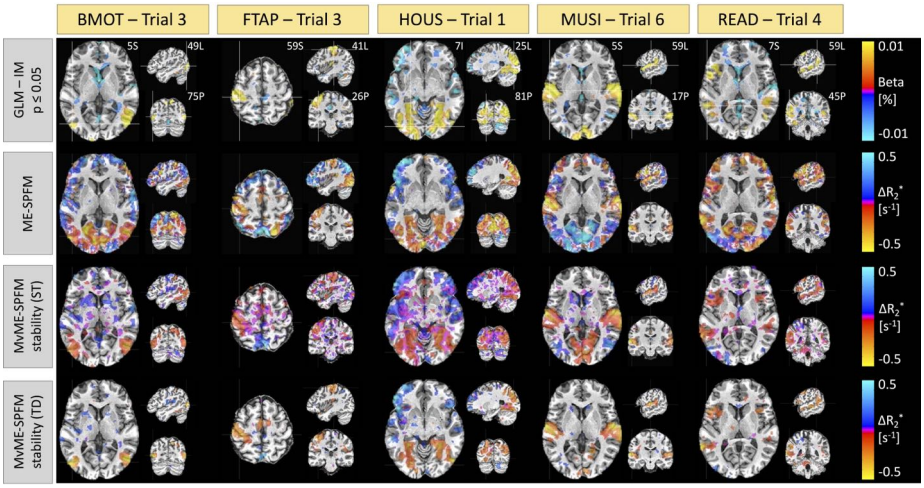

Figure 4.7: Comparison of single-trial activation maps obtained with a GLM (row 1) thresholded at $p \leq 0.05$, the original ME-SPFM formulation with a fixed selection of λ (row 2), the novel MvME-SPFM technique with stability selection, $\rho = 0.5$ and a static threshold (ST, row 3), and using a time-dependent threshold (TD, row 4). A representative trial is shown for each task. All the maps correspond to the same subject and run shown in Figure 4.4.

single-trial ΔR_2^* activity maps and the reference GLM activity maps ($p \leq 0.05$) demonstrates only a slight improvement over the original ME-SPFM formulation when employing an ST thresholding approach with the novel MvME-SPFM technique. In contrast, the Dice coefficients obtained with TD thresholding show a very notable increase of nearly 50% in the median of the distribution of Dice coefficients compared with the original ME-SPFM approach. Similarly, the sensitivity and specificity distributions of ST thresholding demonstrate a slight improvement with respect to the original ME-SPFM formulation. On the other hand, the use of TD thresholding offers nearly perfect specificity ($\geq 95\%$) at the cost of reduced sensitivity across all experimental conditions. Hence, increasing the specificity of the ΔR_2^* maps is more beneficial for increasing the concordance with the GLM maps than increasing the sensitivity.

The receiver operating characteristic (ROC) values in Figure 4.9 corroborate the previous observations regardless of the value of ρ used in the MvME-SPFM method. The estimates obtained with the ST threshold reveal an overall higher sensitivity and a slightly higher specificity compared to the original ME-SPFM technique. In contrast, the ROC values for the TD thresholding approach show a clear improvement in specificity but lower sensitivity. These findings are in line with the results shown in Figure 4.3, Figure 4.7 and Figure 4.5, as the Dice and ROC values certify that the use of stability selection yields robust activation maps regardless of the selection of the spatial regularization term

ρ and obviating the need to choose the temporal regularization parameter λ. An interactive version of Figure 4.8 and Figure 4.9 is available in the following GitHub repository: https://github.com/eurunuela/MvMEPFM_figures.

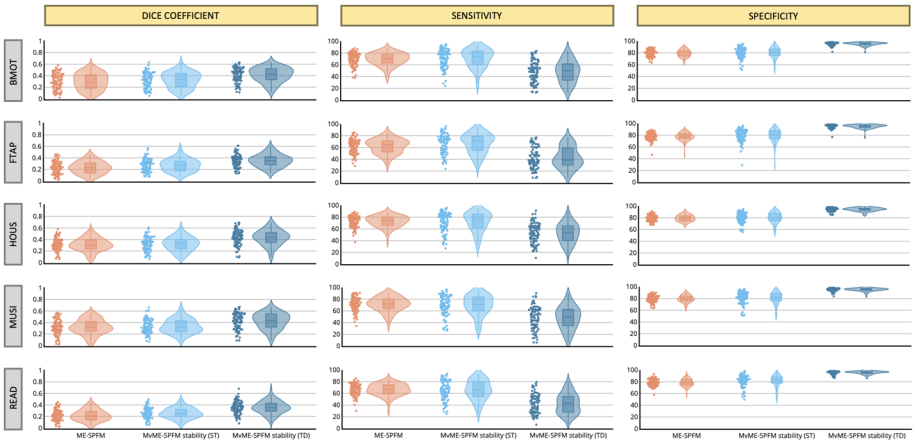

Figure 4.8: Dice coefficient (i.e., spatial overlap), sensitivity and specificity coefficients of the single-trial activation maps for each of the experimental conditions obtained with ME-SPFM, MvME-SPFM with stability selection and a static thresholding approach (ST), and MvME-SPFM with stability selection and a time-dependent thresholding approach (TD). These metrics were obtained with a selection of $\rho = 0.5$. Reference activation maps were obtained with a single trial GLM analysis and thresholded at uncorrected $p \leq 0.05$. The density plot shows the shape of the distribution of the Dice coefficients, and the box plot depicts the median with a solid line, with each box spanning from quartile 1 to quartile 3. The whiskers extend to 1.5 times the interquartile range.

These findings are further corroborated by the average activity maps across trials, sessions and subjects for each condition in the task shown in Figure 4.10. The maps obtained with the novel MvME-SPFM technique with stability selection and TD thresholding show a high resemblance to the average of the single-trial GLM activation maps (3dDeconvolve with independently-modulation (IM) option in AFNI, thresholded at $p < 0.05$). Interestingly, the MvME-SPFM approach seems to offer improved detection of neuronal-related activity than the GLM approach for certain conditions, for example, in the voxels of the left inferior frontal gyrus (i.e., Broca's region) and left posterior superior temporal gyrus (i.e., Wernicke's region) for the reading condition, where the GLM maps show a smaller cluster of activation.

Figure 4.11 shows the comparison with the HemoLearn algorithm (Cherkaoui et al., 2021). Panels A and B illustrate the difficulty of selecting a suitable number of components K for Hemolearn since the number of components giving the highest correlation to the session-level GLM was different for each condition. The timecourses of the estimated activity-inducing signal in C show that Hemolearn

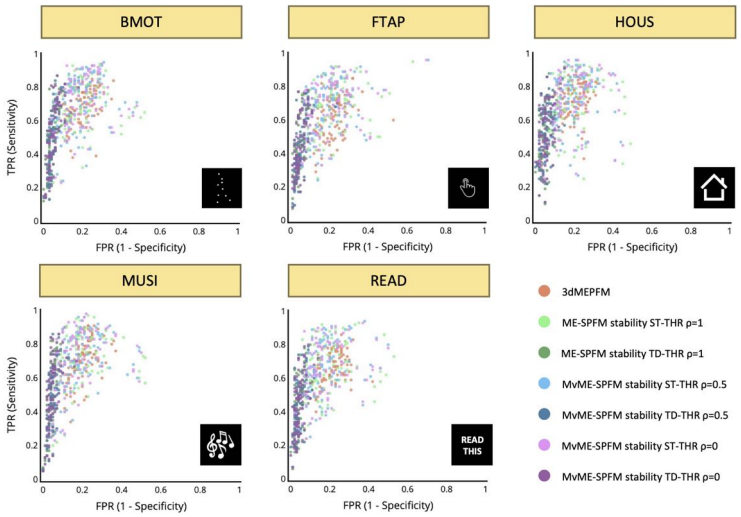

Figure 4.9: Receiver operating characteristic (ROC) values with the sensitivity and specificity of each single trial's activation map for all conditions and the reference map obtained with a single-trial GLM. Different colors are used for the different analyses: the original ME-SPFM, and the novel MvME-SPFM approach using stability selection with the spatial regularization parameter set to $\rho = 0.5$ and the two different thresholding methods: static (ST) and time-dependent (TD). In each analysis each dot represents a single trial, depicting all trials across all datasets, and the centroids across all the single trials are also shown for the three analyses.

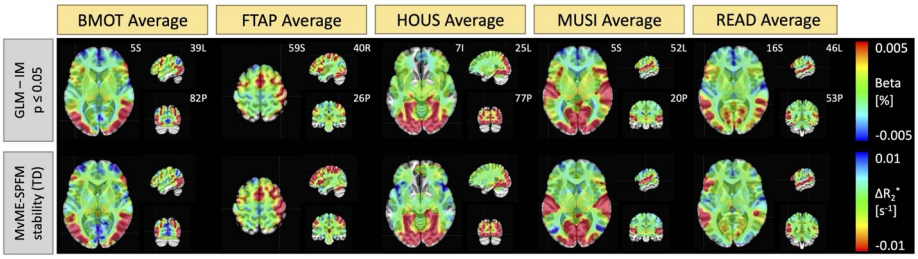

Figure 4.10: Average activity maps across trials, sessions and subjects for each condition in the task obtained with (top) a single-trial GLM analysis (independently modulated, IM) and the proposed MvME-SPFM algorithm with stability selection, $\rho = 0.5$ and a time-dependent threshold (TD).

struggles to capture the timings of the task in all conditions, while the MvME-SPFM technique with stability selection and TD thresholding is able to capture the expected activity-inducing signal with barely missing any of the events in

the task. These observations are also visible in the activity maps shown in D, where the activity maps obtained with the MvME-SPFM technique with stability selection and TD thresholding are comparable to those obtained with the GLM (see Figure 4.7), while the maps of the Hemolearn components that had the highest correlation with the session-level GLM maps do not show a clear activity pattern related to the task.

4.6 Discussion and Conclusion

The proposed whole-brain (i.e., multivariate) formulation for hemodynamic deconvolution of multi-echo fMRI data with the use of stability selection achieved a closer agreement with the activation maps obtained with a single-trial GLM analysis than the original ME-SPFM method (Caballero-Gaudes et al., 2019) and other state-of-the-art multivariate deconvolution method like Hemolearn (Cherkaoui et al., 2021), while obviating the need to select the temporal regularization parameter λ (as shown in Figure 4.7 and Figure 4.11).

Furthermore, this chapter demonstrated that the performance of the proposed method was robust not only for the analysis of multi-echo data, but also when analyzing single-echo data, as demonstrated in Figure 4.6. In addition, our results illustrated that the stability selection procedure also offers robustness against the choice of the spatial regularization parameter ρ, as the AUC maps for different selections of ρ were practically identical, as shown in Figure 4.3. Hence, although stability selection could be employed with a double selection of the regularization parameters λ and ρ, this can be avoided for computational reasons with little influence in the results. In any case, extending the proposed stability selection technique to other formulations of the hemodynamic deconvolution problem, such as the voxel-wise (i.e., univariate) single-echo (Gaudes et al., 2013; Uruñuela et al., 2020) described in Chapter 2 and Chapter 3, univariate multi-echo (Caballero-Gaudes et al., 2019), or low-rank and sparse formulations (Cherkaoui et al., 2021; Uruñuela et al., 2021) described in Chapter 6, is relatively straightforward.

One of the most interesting features of the proposed stability selection procedure is the estimation of the area under the curve (AUC) measure, which provides a new perspective for exploring fMRI data: a 4D movie with the probability of each voxel and time point containing a neuronal-related event. Therefore, the AUC time-series and maps provide complementary information to the estimates of ΔR_2^*, and serve as a reliability measure. Even though the AUC measures were employed here to produce the final estimates of the activity-inducing signal, they could also be exploited on their own. For instance, they could be exploited to constrain functional connectivity analysis (Faskowitz et al., 2020; Tagliazucchi et al., 2016) to voxels and instants with a high probability of containing a neuronal-related event. Furthermore, the stability selection and the AUC metric can also be interpreted from a machine learning perspective, where the outputs from a collection of lasso learners are combined with an ensemble regression approach. Alternatively, the stability selection procedure could also

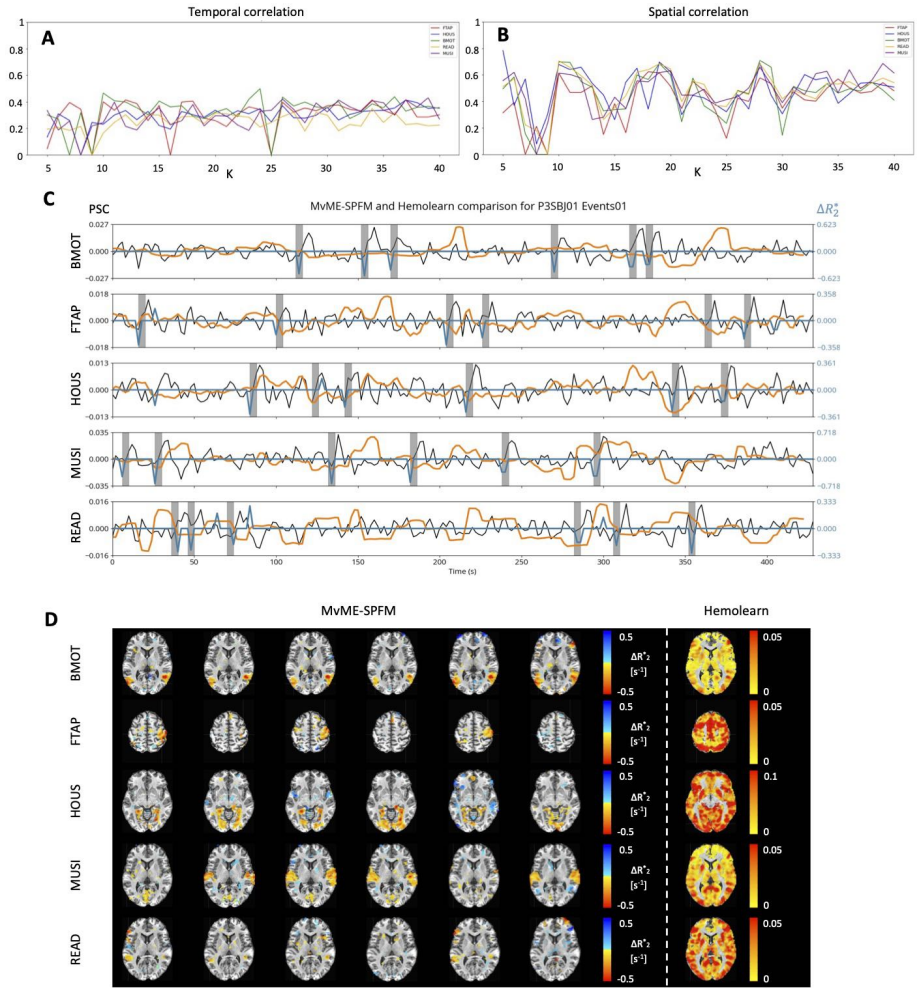

Figure 4.11: Comparison of the activity maps obtained with the novel MvME-SPFM technique with stability selection, $\rho = 0.5$ and a time-dependent threshold (TD) with Hemolearn. Session 1 of subject P3SBJ01 was used for this comparison. A and B show the highest correlation to the session-level GLM regressor and activity map respectively obtained by Hemolearn for each condition in the task (FTAP in red, HOUS in dark blue, BMOT in green, READ in yellow, and MUSI in purple) and the range of K components explored ($K = 5 \ldots 40$). C illustrates the timecourses of MvME-SPFM (blue, ΔR_2^*) and Hemolearn-estimated (orange, percent signal change) activity-inducing signal in a representative voxel for each condition.

be linked to Bayesian approaches where the prior is given by the range of values of the regularization term λ and the total posterior probability of the neuronal

Figure 4.11: The gray bars show the onset and duration of each trial within a condition, and the measured fMRI signal is shown in percent signal change (PSC) in black. D compares the trial-level activity maps estimated with MvME-SPFM with the Hemolearn activity map with the highest correlation with the session-level GLM map across the explored range of K components. Every row in C and D correspond to the different conditions in the task.

event is calculated as the integration of the stability paths, i.e., the AUC (see discussion in Meinshausen & Bühlmann, 2010).

Although the estimation of the AUC eliminates the need to select the spatial and temporal regularization parameters λ and ρ, it requires the use of a thresholding approach given the nature of the AUC measure, which cannot be equal to zero by definition. This chapter adopted two data-driven thresholding strategies, static (ST) and time-dependent (TD), based on the AUC values of a region where no BOLD signal changes related to neuronal activity are assumed to occur (e.g., deep white matter voxels). The use of a static AUC thresholding approach yielded higher sensitivity than the original ME-SPFM method (Caballero-Gaudes et al., 2019) while maintaining the specificity as demonstrated in Figure 4.9. Notably, this improvement was seen in all trials with the exception of one outlier run, regardless of the choice of the spatial regularization parameter ρ. Nevertheless, the use of a time-dependent thresholding approach may be even justified by the increased specificity and nearly perfect retrieval of the activity-inducing signal (row 3 in Figure 4.5) when motion- and respiration-related artifacts are visible in the data (see arrows in Figure 4.4). However, the application of the time-dependent threshold may reduce sensitivity at the single-trial level in some cases. Hence, the results shown in Figure 4.9 encourage the use of the static thresholding approach as an exploratory step before employing the time-dependent threshold. Other thresholding criteria could involve the comparison of AUC values obtained from surrogate (null) data (Liégeois et al., 2021) with the AUC values obtained with the original data.

Furthermore, the extension of the original ME-SPFM algorithm from a voxel-wise to a whole-brain (i.e., multivariate) regularized problem paves the way for more refined formulations that exploit the spatial characteristics and information available in fMRI and complementary imaging data into the spatial regularization term in order to improve the estimation of ΔR_2^*. For instance, the spatial regularization could be constrained within brain regions delineated by commonly used parcellations (e.g., the Schaefer-Yeo atlas) (Karahanoğlu et al., 2013) or within neighbouring gray matter voxels (Farouj et al., 2017). Other mixed-norm regularization terms could also be investigated in future work to account for alternative models of spatially varying noise SNR, e.g., the octagonal shrinkage and clustering algorithm for regression (OSCAR) (Gueddari et al., 2021), which employs a ℓ_1 plus a pairwise ℓ_∞-norm between voxels instead of a global ℓ_2-norm to account for spatially varying SNR across voxels. Moreover, the

multivariate formulation could exploit complimentary multimodal information such as structural connectivity from diffusion-based MRI data (Bolton et al., 2019). In addition, the proposed formulation can be easily adapted to model the changes in neuronal activity in terms of its innovations, which can be more appropiate to capture sustained BOLD events as described in Chapter 2 (see also (Uruñuela et al., 2023)).

One limitation of the proposed MvME-SPFM technique is the assumption of a particular shape of the hemodynamic response to construct the HRF matrix for deconvolution in Eq. (4.3). The proposed model does not account for the variability in the temporal characteristics of the HRF across the brain, which originates from differences in stimulus intensity and patterns, short inter-event intervals, or differences in the HRF shape between resting-state and task-based paradigms (Chen et al., 2021; Polimeni & Lewis, 2021; Sadaghiani et al., 2009; Yeşilyurt et al., 2008). To resolve this issue, given that the performance of MvME-SPFM is not time-locked to the trials, the current formulation could be extended to account for variability in the onset of the activity-inducing signal, as well as to introduce flexibility in the model, by employing multiple basis functions (Gaudes et al., 2012). Furthermore, our method could be easily employed independently within parcels of any commonly-used atlases with a pre-estimated, localized HRF. Finally, the computational demands involved in the stability selection procedure, which solves the regularization problem in Eq. (4.5) for a range of λ values on a number of subsampled surrogate datasets, are higher than solving the regularization path and finding an adequate solution via model selection criteria as in ME-SPFM (Caballero-Gaudes et al., 2019). At the moment of writing, the method took approximately 10 hours to run on a high-performance computing cluster parallelizing the stability selection procedure so that each surrogate dataset was processed in a different core.

4.7 Code and Data Availability

The code and materials used to generate the figures in this work can be found in the following GitHub repository: https://github.com/eurunuela/MvMEPFM_figures.

The Python package is available as part of splora in the following GitHub repository: https://github.com/ParadigmFreeMapping/splora.

Chapter

5

Multi-Subject Paradigm Free Mapping

This chapter is in preparation to be published and will be available as an interactive website at https://github.com/eurunuela/msPFM_paper.

Functional magnetic resonance imaging (fMRI) during naturalistic paradigms has become increasingly popular in cognitive neuroscience as a way to investigate cognition and behaviour within more ecologically valid contexts. However, current methods for fMRI data analysis in naturalistic paradigms are limited in how they can characterize the precise spatio-temporal dynamics of neural activity evoked by naturalistic stimuli. On the one hand, model-based approaches, such as traditional general linear model (GLM) analyses and encoding models, make strong assumptions about what features of the stimulus might be important and when they ocurr, which limits an accurate interpretation of the resulting maps of blood oxygenation-level dependent (BOLD) activity. On the other hand, model-free approaches, such as inter-subject correlation and event segmentation algorithms, make no assumptions about stimulus features, but necessitate combining information across multiple time points, making it difficult to trace the recovered signal back to particular moments in the stimulus. To address these limitations, this chapter introduces a novel method for fMRI data analysis in naturalistic paradigms, called multi-subject Paradigm Free Mapping (msPFM). Unlike existing approaches, msPFM estimates shared and individualized neuronal-related brain responses at the finest spatial and temporal resolution available in fMRI data. The msPFM algorithm implements a multivariate regularized least-squares algorithm for hemodynamic deconvolution of fMRI data that estimates the neuronal-related BOLD activity without relying on the temporal structure of the stimuli. Its performance is demonstrated through simulations and experimental movie-watching datasets, showcasing its

ability to detect moment-to-moment neuronal-related activity patterns associated with low- and mid-level features of the movie. Furthermore, msPFM reveals a spatial cortical gradient of BOLD activity that aligns with previous findings obtained with event segmentation methods. It also provides an extra temporal dimension to the maps obtained with intersubject correlation analyses. In summary, msPFM offers a powerful tool for exploring the spatio-temporal dynamics of the rich patterns of neural activity evoked by naturalistic fMRI paradigms. Researchers can access the open-source Python package for msPFM at https://github.com/ParadigmFreeMapping/msPFM.

5.1 Introduction

The use of movies and stories in the so-called naturalistic paradigms combined with functional magnetic resonance imaging (fMRI) has increased considerably in recent years as a replacement for task and resting state experimental paradigms to study human cognition and mapping brain function (Finn, 2021). It owes its popularity not only to its similarity with freeform cognition and unconstrained behaviour, but also to the fact that individual differences in functional connectivity are more stable (Vanderwal et al., 2017; Wang et al., 2017) and offer improved predictions of behavior (Finn & Bandettini, 2021). Furthermore, naturalistic paradigms have emerged as valuable tools for investigating various neurological and developmental disorders. For example, studies utilizing movie-watching have revealed distinct functional network configurations among individuals with autism, highlighting their idiosyncratic patterns (Bolton, Freitas, et al., 2020; Bolton et al., 2018). Additionally, research has demonstrated a reduction in intersubject correlations within attention-related brain areas among individuals with attention deficit hyperactivity disorder (ADHD), indicating a desynchronization of brain activity in this diseased population (Salmi et al., 2020). Movie-watching fMRI has also been utilized to investigate the memory representations of healthy individuals during the retrieval of movie scenes, which revealed a shared spatial organization among subjects exposed to identical naturalistic stimuli (Chen et al., 2017). Finally, naturalistic paradigms have played a crucial role in demonstrating the meaningful structure and reliability of movie scene segmentation across subjects (Raccah et al., 2022; Speer et al., 2009; Zacks et al., 2009). Notably, studies have revealed that variations in these patterns of event segmentation are associated with age, cognitive abilities (Bailey et al., 2017; Jafarpour et al., 2022; Zacks et al., 2006), and clinical conditions (Zalla et al., 2004).

Several techniques have been employed to analyze naturalistic fMRI data. One such technique is intersubject correlation (ISC), which involves calculating the synchrony among the voxelwise signals of multiple subjects throughout the data acquisition process (Chen et al., 2020; Hasson et al., 2004; Jangraw et al., 2023; Nastase et al., 2019; Wilson et al., 2007). An extension to this approach incorporates the temporal dimension with the sliding window ISC method (averaging ISC over a temporal window of a few volumes) or the

instantaneous phase synchronization (Glerean et al., 2012), These approaches generate voxelwise measures of moment-to-moment synchrony (Nummenmaa et al., 2012), which can then be contrasted against stimulus models, offering an intermediate approach that combines aspects of stimulus-dependent and stimulus-free analyses. However, the degree of synchrony is calculated during the entire acquisition or within the win dow duration, i.e. spanning over multiple TRs. Intra and inter-subject correlation of functional connectivity (Vanderwal et al., 2017) and the intersubject correlation of dynamic connectivity (Di & Biswal, 2020) offer additional approaches for investigating naturalistic fMRI data that examine the interplay between brain regions and the consistent patterns of connectivity across individuals. Alternatively, the detection of event segmentation boundaries has also been a major focus in the development of data analysis techniques for naturalistic paradigms. Currently, two main state-of-the-art techniques are mostly used to blindly identify temporal boundaries within a sequence of events without prior knowledge: hidden Markov models (HMM) (Baldassano et al., 2017) and greedy state boundary search (GSBS) (Geerligs et al., 2021). Finally, another strategy involves constructing encoding models that utilize representations extracted from neural network language models to forecast BOLD responses to natural language stimuli (Jain et al., 2020), as well as the use of deep neural network models using human visual areas to learn non-hierarchical representations (St-Yves et al., 2023).

However, the aforementioned methods have several limitations. First, ISC reveals widespread synchronization across individuals without specific associations with events of particular cognitive functions throughout a scan. Although sliding-window ISC can enhance temporal resolution for characterizing such events, it should be noted that this technique does not offer the utmost level of temporal precision. Second, the HMM and GSBS operate at the ROI level, and hence provide a low spatial resolution that does not allow the discovery of precise, individualized spatial patterns. More importantly, all these techniques summarize the within-region fMRI signal into states or connectivity patterns that are sustained over several TRs, and therefore, they do not capture temporal dynamics at the fastest temporal resolution available in the data, i.e., at the TR level. In summary, these methods rule out the precision mapping of idiosyncratic spatio-temporal patterns evoked by naturalistic stimuli. One possible approach to address this issue is to employ data-driven techniques that operate at the TR and voxel level. Hemodynamic deconvolution is a method that falls under this category (Gaudes et al., 2010; Karahanoğlu et al., 2013; Uruñuela et al., 2023). It has the ability to estimate neuronal-related activity without relying on the temporal information of the experimental paradigm, and as a result, it proves to be particularly valuable when dealing with naturalistic stimuli.

Here, we introduce a novel method for the analysis of naturalistic fMRI data called multi-subject Paradigm Free Mapping (msPFM) that overcomes the limitations mentioned above. msPFM solves a multivariate regularized least-squares problem for hemodynamic deconvolution of the fMRI signal across subjects to blindly estimate the shared and individualized neuronal-related activity without requiring any temporal model of the stimulus. Unlike previous

methods for naturalistic fMRI data, this formulation allows the msPFM algorithm to operate at the fastest temporal resolution (TR), and the highest spatial resolution (voxel). We introduce msPFM and evaluate the method in simulations and two experimental datasets with movie-watching paradigms to showcase its potential to elucidate moment-to-moment spatio-temporal neural activity patterns evoked by naturalistic paradigms and assess their ability to map low- and mid-level movie features. The datasets, Iteration and Sherlock, differed in terms of the studied low- and mid-level features, subject counts (43 vs 16), and data type (multi-echo vs single-echo).

5.2 Multi-subject Paradigm Free Mapping

5.2.1 Signal Model

In gradient-echo fMRI acquisitions, the time series of a voxel v normalized to signal percentage change can be approximated with the following linear relationship with the echo time (TE) (Kundu et al., 2017):

$$y(v, TE_k, t) \approx \Delta\rho(v, t) - TE_k \Delta R_2^*(v, t) + \epsilon(v, t), \tag{5.1}$$

where $\Delta R_2^*(v, t)$ represents BOLD-like changes, $\Delta\rho(v, t)$ describes changes in the net magnetization, and $\epsilon(v, t)$ is a random noise term. Changes in the net magnetization $\Delta\rho(v, t)$ are mostly due to motion- and hardware-related artefacts, which can be reduced in preprocessing by using denoising tools such as ICA-AROMA (Pruim, Mennes, van Rooij, et al., 2015) for single-echo acquisitions or ME-ICA (Kundu et al., 2012, 2017) for multi-echo acquisitions. Hereinafter, the term $\Delta\rho(v, t)$ can be neglected for a simpler notation.

Following the same notation described in previous chapters, let us also consider that changes in $R_2^*(v, t)$ that are related to neuronal activity will follow a linear time-invariant model that is defined as the convolution of the activity-inducing signal $s(v, t)$ with the hemodynamic response $h(v, t)$ (Boynton et al., 1996; Glover, 1999). Considering the signal is sampled at every TR, the previous signal model can be written in discrete matrix form as $\mathbf{y}_v = -TE_k \mathbf{H}\mathbf{s}_v + \boldsymbol{\epsilon}_v$, where $\mathbf{H} \in \mathbb{R}^{N \times N}$ is the Toeplitz convolution matrix with shifted versions of the hemodynamic response function (HRF) (Gaudes et al., 2013; Gitelman et al., 2003), $\mathbf{y}_v$, $\mathbf{s}_v$, $\boldsymbol{\epsilon}_v \in \mathbb{R}^N$, and N is the number of scans (i.e. TRs).

Without loss of generalization, we assume that the HRF model is the same across voxels. As described in Chapter 2, the activity-inducing signal $\mathbf{s}_v$ can also be written in terms of the innovation signal $\mathbf{u}_v$ as $\mathbf{s}_v = \mathbf{L}\mathbf{u}_v$, where the matrix $\mathbf{L} \in \mathbb{R}^{N \times N}$ is the first-order integration operator (Cherkaoui et al., 2019; Uruñuela et al., 2020, 2023). Hence, the signal model for a given TE_k (i.e. single-echo) can be written as follows:

$$\mathbf{y}_{v,k} = -TE_k \mathbf{H}\mathbf{L}\mathbf{u}_v + \boldsymbol{\epsilon}_v, \tag{5.2}$$

which can be extended for a multi-echo acquisition by concatenating the time series of each echo ($k = 1, \ldots, K$) in a single vector:

$$\begin{bmatrix} \mathbf{y}_{v,1} \\ \vdots \\ \mathbf{y}_{v,K} \end{bmatrix} = \begin{bmatrix} -\mathrm{TE}_1\mathbf{HL} \\ \vdots \\ -\mathrm{TE}_K\mathbf{HL} \end{bmatrix} \mathbf{u}_v + \begin{bmatrix} \boldsymbol{\epsilon}_{v,1} \\ \vdots \\ \boldsymbol{\epsilon}_{v,K} \end{bmatrix}. \tag{5.3}$$

Assuming that data from multiple subjects ($s = 1, \ldots, S$) are aligned to a common space, the multivariate multi-echo formulation introduced by (Uruñuela et al., 2022) (see Chapter 4) is adapted to simultaneously estimate shared and individualized voxelwise activity-inducing signal in multiple subjects. This method, referred to as multi-subject Paradigm Free Mapping (msPFM), therefore proposes to concatenate the voxelwise fMRI signal of multiple subjects into $\mathbf{Y}_v \in \mathbb{R}^{K \cdot NS}$. Hence, for a given voxel v or ROI the signal model of msPFM can be written as follows:

$$\begin{bmatrix} \mathbf{y}_{v,1,1} & \cdots & \mathbf{y}_{v,S,1} \\ \vdots & \ddots & \vdots \\ \mathbf{y}_{v,1,K} & \cdots & \mathbf{y}_{v,S,K} \end{bmatrix} = \begin{bmatrix} -\mathrm{TE}_1\mathbf{HL} \\ \vdots \\ -\mathrm{TE}_K\mathbf{HL} \end{bmatrix} \begin{bmatrix} \mathbf{u}_{v,1} & \ldots & \mathbf{u}_{v,S} \end{bmatrix} + \begin{bmatrix} \boldsymbol{\epsilon}_{v,1,1} & \ldots & \boldsymbol{\epsilon}_{v,S,1} \\ \vdots & \ddots & \vdots \\ \boldsymbol{\epsilon}_{v,1,K} & \ldots & \boldsymbol{\epsilon}_{v,S,K} \end{bmatrix}, \tag{5.4}$$

which can be simplified as $\overline{\mathbf{Y}} = \overline{\mathbf{H}}\mathbf{U}_v + \mathbf{E}$.

5.2.2 Inverse Problem

Assuming the noise follows a Normal distribution, the innovation signals for all subjects can be estimated by solving a multivariate ordinary least squares problem. However, due to high collinearity of the Toeplitz convolution matrix $\mathbf{H}$, the estimates will exhibit a large variance in temporal space. Furthermore, in this multi-subject setting where subjects are exposed to the same naturalistic stimuli, certain amount of synchronization of the neuronal activity across subjects can be presumed at given instances. Based on these priors, the estimation of the innovation signal $\mathbf{U}_v$ can be solved with the following regularized multivariate least squares problem:

$$\mathbf{U}_v = \arg\min_{\mathbf{U}_v} \|\overline{\mathbf{Y}}_v - \overline{\mathbf{H}}\mathbf{U}_v\|_2^2 + \rho\|\mathbf{D}_v\mathbf{U}_v\|_1 + (1-\rho)\|\mathbf{D}_v\mathbf{U}_v\|_{2,1}. \tag{5.5}$$

Here, the regularization term $\|\mathbf{D}_v\mathbf{U}_v\|_1$ promotes sparsity of the innovation signals across time and subjects, treating each estimate independently. On the other hand, the mixed-norm regularization term $\|\mathbf{D}_v\mathbf{U}_v\|_{2,1}$ promotes the co-activation of the innovation signals across all subjects simultaneously (Gramfort et al., 2011; Uruñuela et al., 2022). The parameter ρ controls the balance between the two regularization terms, and $\mathbf{D}_v = \mathrm{diag}(\lambda_{v,1}, \ldots, \lambda_{v,S}) \in \mathbb{R}^{S \times S}$ is a diagonal matrix with subject-specific and region-specific regularization parameters $\lambda_{v,s}$, which allows adjusting to the level of noise in each subject's data and region.

As it has been discussed in previous chapters, the selection of the regularization parameters $\lambda_{v,s}$ is critical to obtain accurate estimates of the

innovation signal $\mathbf{U}$. Here, for each region, a subject-specific regularization parameter $\lambda_{v,s} = 30\sigma_s$ was chosen for each subject s, where σ_s is a pre-estimated noise level of the data fit and the pre-estimated noise level is derived from the median absolute deviation of fine-scale wavelet coefficients (Daubechies, order 3) of each subject. In addition, the parameter ρ was set to 0.8 to balance the sparsity of the innovation signals (i.e. capturing the idiosyncrasies of each subject and moment) and the grouping of synchronous co-activations across subjects (i.e. shared neuronal responses across subject at a given time). Both the λ and ρ parameters were set based on simulations.

Relevantly, notice that msPFM estimates a 4D dataset of the innovation signal $\mathbf{U}_v \in \mathbb{R}^{T \times S}$ for every subject. That is, unlike a GLM method that calculates activation maps for each onset in the paradigm, msPFM estimates activity maps for each TR. In fact, the innovation signal estimated by msPFM can easily be converted into the activity-inducing signal with an integration operator, and convolved with the HRF to obtain a denoised reconstruction of the data.

5.3 Methods

5.3.1 Simulations

We generated a voxel time-series for subjects divided into three groups (groups 1-3) with similar, yet different BOLD patterns to simulate group responses and idiosyncrasies as shown in Figure 5.1 and Figure 5.2. We also generated two time-series with no BOLD responses (group 4) as a time-series were msPFM should not find any activity. The time-series included Gaussian noise to simulate the effect of noise in fMRI data, and were normalized to percent signal change.

Two scenarios were simulated to study the effect of how the relative number of subjects in each group affects the estimation of the innovation signal in the other subjects: (1) the groups contained a different ratio of subjects per group: group 1 contained 35 subjects, group 2 contained 15 subjects, and group 3 contained 5 subjects (see Figure 5.1); and (2) the same ratio of subjects per group is preserved: groups 1, 2 and 3 contain the same amount of subjects (20) as shown in Figure 5.2.

The data for both scenarios were analyzed with the msPFM algorithm, where λ_s for each subject was carefully chosen as $\lambda_s = 30\sigma_s$, as described in detail in Section 5.2.2 The parameter ρ was set to 0.8. A Jupyter Notebook with all the simulated scenarios is available at https://github.com/eurunuela/msPFM_paper.

5.3.2 Experimental Data: MR Acquisition and Preprocessing

The performance of the proposed msPFM approach is illustrated in two naturalistic movie-watching datasets acquired with different MRI parameters and with different cognitive characteristics: The Iteration dataset and the Sherlock dataset. These datasets were chosen due to their diffferent characteristics, as they would allow to validate the proposed msPFM method in a broader manner.

In particular, Iteration offers 43 subjects and multi-echo data as opposed to Sherlock's 16 subjects and single-echo data. However, Sherlock has been widely used by the naturalistic fMRI community and thus is well known. In addition, Sherlock contains a higher number of characters compared with the single character in Iteration (a second one appears at the very end), which implies that there is a higher amount of speech present in Sherlock. Finally, the movies are also different in terms of their visuals, with Iteration mostly having a bright background, while Sherlock contains more variety in terms of brightness.

The Iteration Dataset

48 participants were used in this study (all native English speakers; 27F, median age 24.5, range = (19,64)) from the National Institutes of Health (NIH). See Sava-Segal et al., 2022 for full procedural detail but in brief: subjects watched four movies (ranging from 7:27-12:27 min each) and subsequently completed a task battery designed to probe their interpretations and reactions to the movie, including the following: (i) a 3-min free recall/appraisal task in which subjects spoke freely about their memories and impressions of the movie, during which their speech was captured with a noise-canceling microphone; (ii) multiple-choice comprehension questions designed to ensure they had been paying attention; (iii) multiple-choice and Likert-style items assessing reactions to various characters and to the movie overall. The movie order was pseudorandomized for each subject such that order was counterbalanced at the group level. From the four movies, we selected the one with the clearest event boundaries or scene cuts.

MRI data were collected using a 3T Siemens Prisma scanner with a 64-channel head coil in 43 subjects (all native English speakers; 27F, median age 24.5, range= (19,64)) while they watched a movie (called Iteration; https://youtu.be/c53fGdK84rc - 12:27 min:sec) at the National Institutes of Health (NIH). All subjects provided informed written consent prior to the start of the study in accordance with experimental procedures approved by the Institutional Review Board of the NIH. Functional images were acquired using a T2*-weighted multiband multi-echo echo-planar imaging (EPI) pulse sequence (TR = 1000 ms, echo times (TE) = [13.6, 31.86, 50.12 ms], flip angle = 60 deg, field of view = 216 x 216 mm, voxel resolution = 3.0 mm isotropic, number of slices = 52 (whole-brain coverage), multiband acceleration factor = 4, with interleaved acquisition). Anatomical images were acquired using a T1-weighted MPRAGE pulse sequence (TR = 2530 ms, TE = 3.30 ms, flip angle = 7 deg, field of view = 256 x 256 mm, voxel resolution = 1.0 mm isotropic). For more details on the data and experimental design, we refer the reader to (Sava-Segal et al., 2022). The raw data are available at https://openneuro.org/datasets/ds004516/.

Following the conversion of the original DICOM images to NIFTI format, we used AFNI (Cox, 1996) to preprocess MRI data. The preprocessing pipeline was generated with AFNI's afni_proc function and edited to use tedana (DuPre et al., 2021; The Tedana Community et al., 2022) for multi-echo independent component analysis (ME-ICA) denoising. Preprocessing included the following steps: despiking, volume realignment, affine coregistration with the subject-

specific anatomical T1-weighted image, nonlinear warping to the MNI152_2009 template, multi-echo ICA decomposition with tedana (DuPre et al., 2021; The Tedana Community et al., 2022), manual classification of ME-ICA components with Rica (Uruñuela, 2023), nuisance regression of the components classified as noise, spatial smoothing with an isotropic gaussian kernel with full-width half-maximum of 4 mm, detrending of up to 6th order polynomials, nuisance regression of each subject's six motion time series, their derivatives, and linear polynomial baselines for each of the functional runs, and normalization to percent signal change. Nuisance regression steps were perfomed together in as single step before normalizing to percent signal change.

Finally, the preprocessed data was then analyzed with a shared response model (Chen, Chen, et al., 2015) in BrainIAK (Kumar et al., 2020) to account for different functional topographies across individuals. First, the model was fit to capture reliable whole-brain responses to the movie across subjects in a lower dimensional feature (or component) space of 50 features. Then, the individual voxelwise time courses for each participant were reconstructed using this model. This procedure serves as an additional denoising step and makes spatiotemporal patterns more consistent across subjects. For computational reasons, all analyses were conducted on the average time-series of the 1000 cortical ROIs of the Schaefer atlas (Schaefer et al., 2017).

The Sherlock Dataset

16 participants from the Sherlock dataset in (Chen et al., 2017) were used for this study. See the original manuscript for full procedural details, but in brief: participants watched a 48-min segment of the BBC television series Sherlock and subsequently verbally recalled the narrative of the show aloud while undergoing fMRI data acquisition. During the recall task participants were instructed to talk for a minimum of 10 minutes, and were allowed to talk for as long as they wished. Experimenters manually ended the scanning run during the recall task based on verbal indication from participants.

MRI data were collected on a 3T Siemens Skyra scanner with a 20-channel head coil. Functional images were acquired using a T2*-weighted echo-planar imaging (EPI) pulse sequence (TR = 1500 ms, TE = 28 ms, flip angle = 64 deg, field of view = 192 × 192 mm^2, in-plane resolution = 3.0 × 3.0 mm^2, slice thickness = 4 mm, number of slices = 27 (whole-brain coverage) with ascending interleaved acquisition). Reported analyses come from a viewing run of 23 minutes. The raw data were available at https://openneuro.org/datasets/ds001132/versions/1.0.0.

The Sherlock data was preprocessed using fMRIPrep (Esteban et al., 2018) with the ICA AROMA option (Pruim, Mennes, van Rooij, et al., 2015). The preprocessing also included spatial smoothing with a gaussian kernel of fwhm=6mm, and voxelwise nuisance regression of the 6 realignment parameters, their squares, their derivatives, and squared derivatives. Finally, we performed a detrending step using a 6th order polynomial to remove slow drifts in the data and normalized the data to percent signal change. Similar to The Iteration dataset, the preprocessed data was then applied the shared response model (Chen, Chen,

et al., 2015) in BrainIAK (Kumar et al., 2020), and analyses were conducted using the average time-series of the 1000 ROIs from the Schaefer atlas (Schaefer et al., 2017).

5.3.3 msPFM and Post Analyses

The msPFM algorithm was employed to estimate the activity-inducing signal for each subject and parcel in the Schaefer 1000-ROI atlas. While msPFM allows for a voxelwise analysis, 1000 ROIs were employed in the analysis for computational efficiency and to address the voxel mismatch issue caused by variations in anatomical structures across subjects. The regularization parameters λ and ρ were selected as described in the simulations section above. The double-gamma SPM HRF was employed via nilearn's first level GLM module.

Population Synchrony

In order to summarize the estimated activity-inducing signal across subjects and to compute metrics that allow us to validate our method, the positive population synchrony was calculated, denoted as PopSync+. This metric quantifies the number of subjects that show positive neuronal-related activity at each TR, i.e., the number of subjects that activate together in response to the stimuli. This metric was calculated for each ROI, thus obtaining a 4D map of positive population synchrony. Note that unlike other methods that have been used in the literature to quantify population synchrony, this method provides a continuous measure at the TR level, rather than a value per temporal window, which allows us to perform a more fine-grained analysis of the population synchrony.

Event Rate and Inter-Subject Correlation

Recent studies have used hidden Markov model (HMM) algorithms to automatically segment brain activity during the experimental acquisition into a sequence of events (a.k.a. event segmentation). A cortical gradient of segmentation rate has been reported such that primary sensory regions segment (transition between events) more quickly, while higher-order association regions segment more slowly (Baldassano et al., 2017; Geerligs et al., 2022); this principle has been related to temporal receptive windows and intrinsic timescales of brain regions (Hasson et al., 2008; Murray et al., 2014). Conversely, these findings suggest a lower occurrence in higher-order association areas that are known to be sensitive to narrative information over longer timescales (Honey et al., 2012; Lerner et al., 2011). The aim here was to determine if msPFM recovers this same cortical gradient. To this end, the msPFM-based event rate was calculated for each parcel, defined as the number of activity-inducing signal events per minute, using the PopSync+ time series. Only events in the activity-inducing signal that elicited a positive BOLD response were used; i.e., events were defined as instances where the activity-inducing signal remained non-zero for at least 5 TRs.

Another known feature of brain activity during naturalistic stimulation is that synchrony of activity time courses across subjects, known as inter-subject correlation (ISC) (Finn et al., 2020; Hasson et al., 2004; Nastase et al., 2019), also roughly follows a sensory-association gradient, such that subjects are more similar to one another (i.e., more synchronized) in primary visual and auditory regions, while activity becomes more idiosyncratic in higher-order regions. An ISC-like metric was defined from msPFM estimates by quantifying the median correlation across subjects of the estimated activity-inducing signal in each parcel. Then, this was compared to a "classic" ISC approach, in which the median correlation across subjects of average raw BOLD signal was quantified in each parcel.

Correlation Between Population Synchrony and Movie Features

In order to confirm that the multi-subject Paradigm Free Mapping method can detect brain activity patterns related to sensory stimuli, the correlation between the population synchrony and the following low-level features was calculated: changes in luminance (i.e., the derivative of the luminance), and changes in audio (i.e., the derivative of the audio envelope). The low-level features were extracted from the movies employing a Python package named `pliers`. More precisely, the `BrightnessExtractor` and `RMSExtractor` functions were used respectively. The correlation between the population synchrony and the low-level features was calculated for each ROI, thus obtaining a 3D dataset of correlation values.

The study of the relationship between the brain activity and the naturalistic stimuli was extended to mid-level features. For instance, naturalistic stimuli where subjects see faces, buildings and hand-related movements evoke brain activity patterns in the fusiform gyrus (Epstein & Kanwisher, 1998; Haxby et al., 2000), the collateral sulcus (Aguirre, Zarahn, & D'Esposito, 1998a; Epstein & Kanwisher, 1998), and the middle postcentral sulcus (Hasson et al., 2004) respectively. In this study, the primary focus was to examine the relationship between population synchrony and speech, the presence of hands, and the presence of faces in the movies. To extract the relevant information, the `ClarifaiAPIVideoExtractor` and `FaceRecognitionFaceLocationsExtractor` functions provided by the `pliers` library we utilized. These functions allow to accurately determine the timing of speech and hand movements, as well as detect the presence of faces in the movie footage.

5.4 Results

5.4.1 Simulations

Benchmarking msPFM on simulated data where the ground truth is known enables to assess whether it can effectively identify shared neuronal responses across subjects, while accommodating individual differences in response to naturalistic stimuli. The two simulation scenarios described in Section 5.3 were employed here.

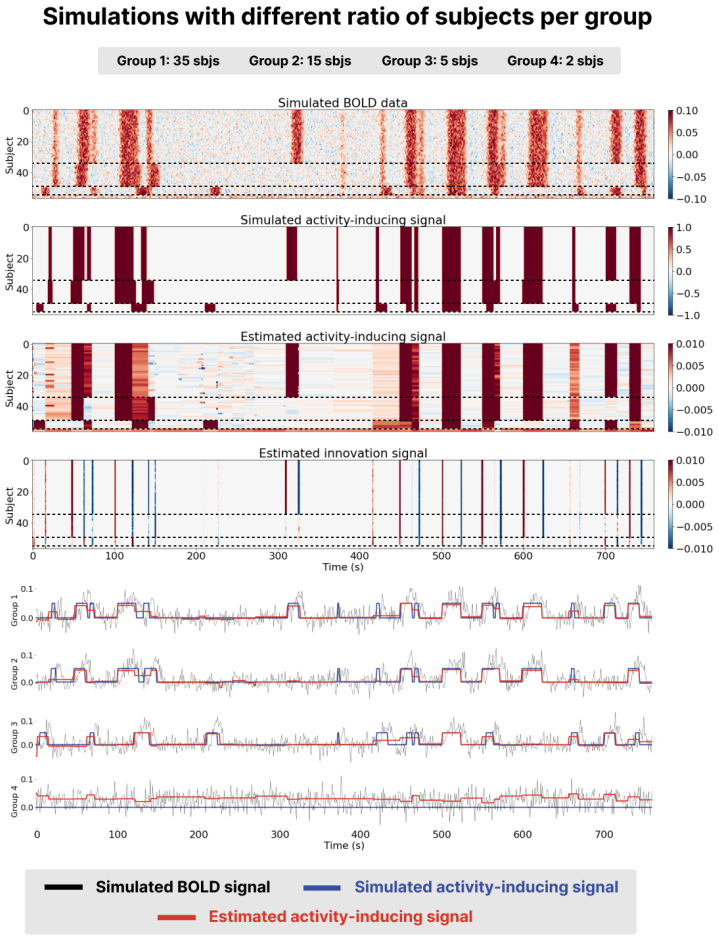

Figure 5.1: Results of running msPFM on simulated data for a ratio of subjects per group of 35, 15, and 5 subjects respectively (groups are separated by a dashed line). The first two heatmaps represent the simulated BOLD and activity-inducing signals for all the simulated subjects. The third and fourth heatmaps represent the estimated activity-inducing and innovation signals for all the simulated subjects. The time courses below depict the simulated BOLD (black) and activity-inducing (blue) signals for a subject in each group and the estimated activity-inducing signal is shown in red.

Visual inspection of the results, as presented in the bottom two heatmaps and the red time courses on Figure 5.1 and Figure 5.2, revealed that the activity-inducing signal estimated by msPFM closely resembled the ground truth. This was true for both scenarios, with an unequal and an equal number of subjects per group. msPFM demonstrates proficiency in estimating prolonged

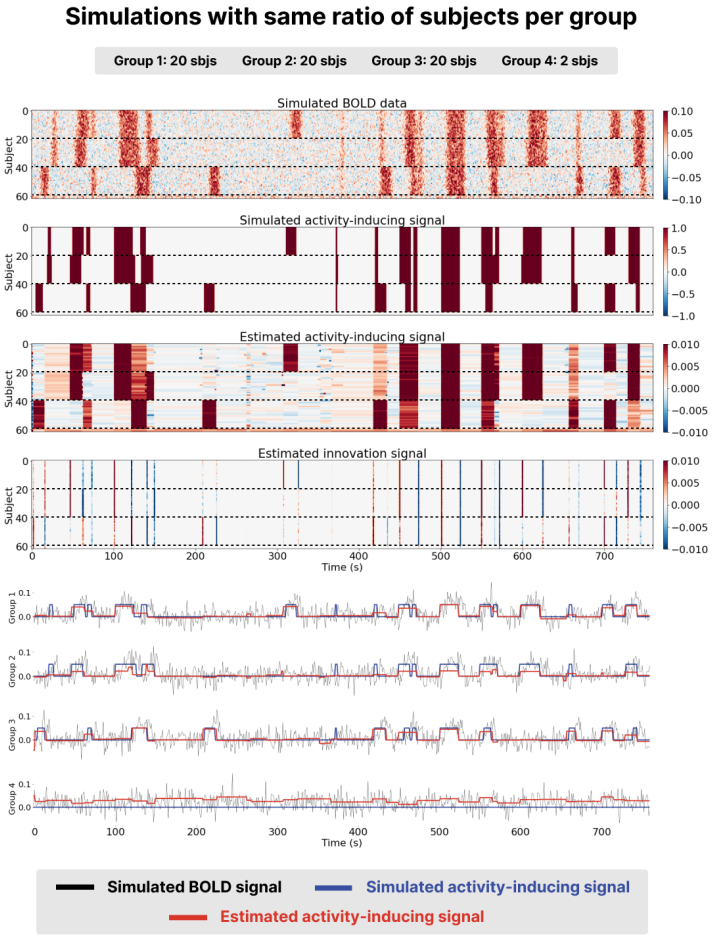

Figure 5.2: Results of running msPFM on simulated data for a ratio of subjects per group of 20, 20, and 20 subjects respectively (groups are separated by a dashed line). The first two heatmaps represent the simulated BOLD and activity-inducing signals for all the simulated subjects. The third and fourth heatmaps represent the estimated activity-inducing and innovation signals for all the simulated subjects. The time courses below depict the simulated BOLD (black) and activity-inducing (blue) signals for a subject in each group and the estimated activity-inducing signal is shown in red.

activity-inducing signals owing to the adequacy of using a model that estimates the innovations, albeit with relatively lower accuracy in capturing shorter, transient activations (we refer the reader to Uruñuela et al., 2023 for an in-depth comparison between models based on the activity-inducing signal –the spike model– or the innovations –the block model). Furthermore, the influence

of the $\ell_{2,1}$-norm (see Equation (5.5)) is conspicuous in the estimation of group four, which lacks simulated activity-signal. In this case, artifacts originating from the estimates of the remaining groups are noticeable, yet these artifacts are relatively minor compared to the estimated activity-inducing signals of the other groups. In addition, the estimated signals for these subjects never return to zero, making it unlikely that they represent any neuronal-related activity and easily identifiable as artefacts. These results are consistent for both scenarios. Hence, based on the simulations, msPFM can pinpoint events that trigger shared and individualized neuronal responses without knowing their timings.

5.4.2 PopSync+ Proof of Concept: Individual and Group Responses to Naturalistic Stimuli

Figure 5.3 and Figure 5.4 show a proof of concept of the PopSync+ metrics and the individual differences in the estimated activity-inducing signal. In fact, the thinner time series in both datasets clearly demonstrate the presence of individual differences in the msPFM estimates of each subject, ehich are consistly observed in the two examples of the left middle temporal gyrus (Brodman area 21) and the V1 visual area. Notably, the PopSync+ time series (depicted in red) displays a pronounced alignment with speech events occurring in the movies (depicted in black). This alignment is particularly prominent in the context of Sherlock, wherein the episode encompasses a notably higher frequency of speech events compared to Iteration. In the case of the latter, msPFM not only captures the speech events but also identifies additional neural dynamics during non-speech periods in this ROI. Figure 5.4 illustrates the synchrony between changes in the luminance of the movies and the PopSync+ in V1. Specifically, a stronger alignment is observed in the case of Iteration, where the changes in luminance exhibit greater extremes compared to the case of Sherlock, which is also reflected in the estimated neuronal response in V1 and hence the PopSync+ measure.

5.4.3 Comparison with Hidden Markov Models and Inter-Subject Correlation

As shown in Figure 5.5, the event rates of both HMM and msPFM for the Iteration dataset show a posterior-to-anterior cortical gradient with a higher frequency of events in posterior, sensory regions and slower rates in anterior, higher-order regions. This gradient is consistent with previous findings in the literature (Sava-Segal et al., 2022). However, the scatter plots in the bottom row illustrate that the difference in event rates between posterior and anterior regions is less pronounced in the case of msPFM as compared with the HMM. This discrepancy can be attributed to the divergent definitions of what an event is for the two methods. The HMM approach identifies states that are expected to persist for a specific number of time points, while the msPFM algorithm captures "punctate" events that can be directly associated with individual moments in the stimulus. As a consequence, the msPFM algorithm is expected to detect

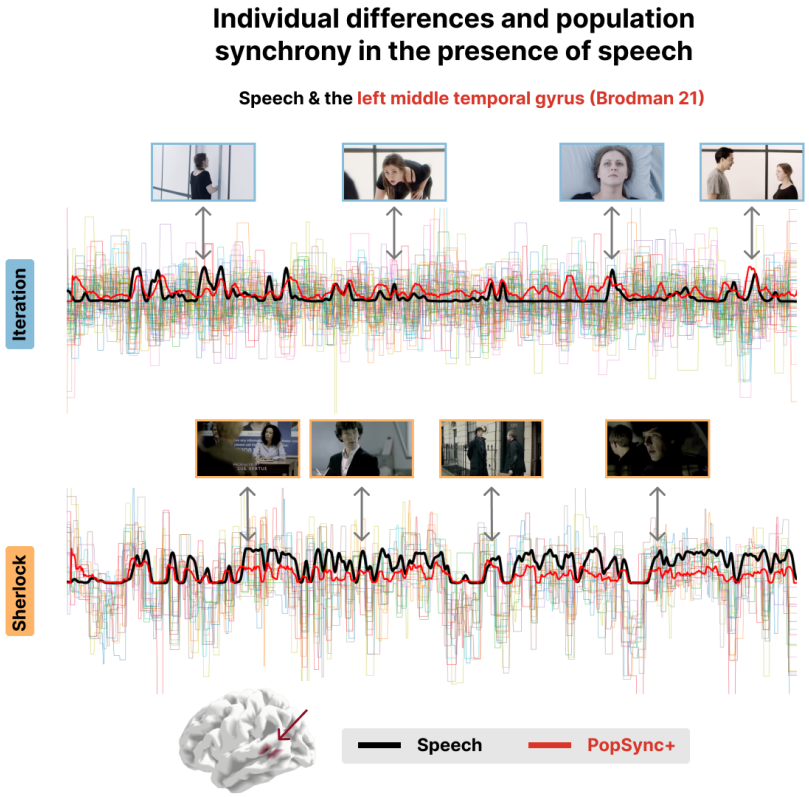

Figure 5.3: Time series of the estimated activity-inducing signal for each subject in the left middle temporal gyrus (Brodman 21 area) for the Iteration (top) and Sherlock (bottom) datasets. The black lines represent the speech events in the movie. The red lines represent the PopSync+, i.e., the sum across subjects of the activity-inducing signal that evokes a positive BOLD response in each parcel. Representative instances of the movies and their respective Speech and PopSync+ TRs are shown for both datasets.

a greater number of events, particularly in anterior regions. In the context of the Sherlock dataset, it is worth noting that the event rate is noticeably lower when calculated with HMM, and consequently, the cortical gradient is less apparent. On the other hand, the msPFM event rate demonstrates values similar to Iteration and reveals a more distinct cortical gradient. When it comes to the inter-subject correlation, the two maps and their values were remarkably similar as shown in Figure 5.6, indicating that msPFM is able to recover this same sensory-association gradient, with subjects being more synchronized with one another lower-order regions, while being more idiosyncratic in higher-order regions.

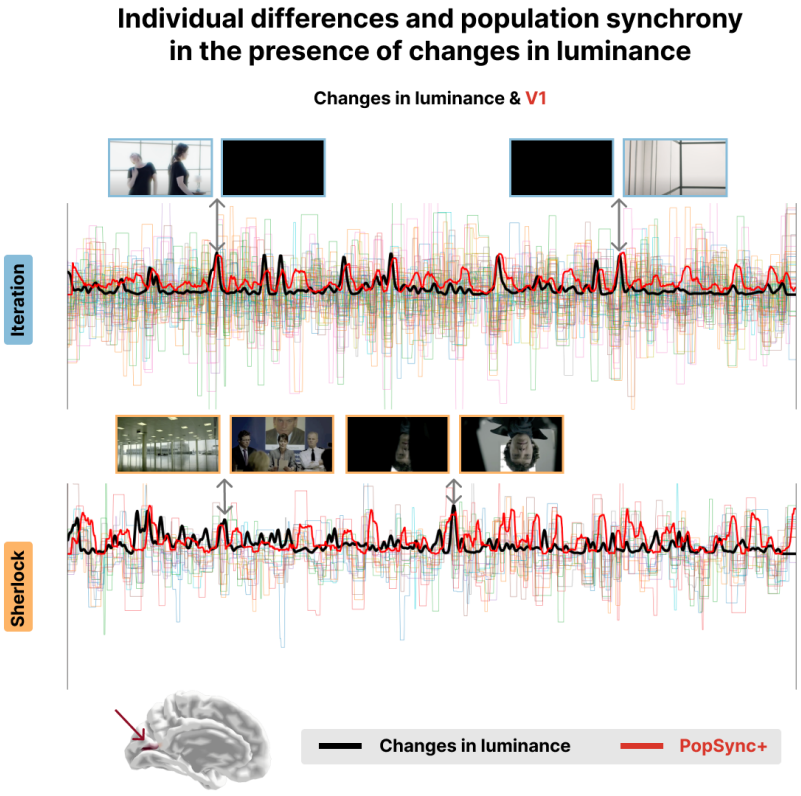

Figure 5.4: Time series of the estimated activity-inducing signal for each subject in V1 for the Iteration (top) and Sherlock (bottom) datasets. The black lines represent the changes in luminance in the movie. The red lines represent the PopSync+, i.e., the sum across subjects of the activity-inducing signal that evokes a positive BOLD response in each parcel. Representative instances of the movies and their changes in luminance and PopSync+ TRs are shown for both datasets.

5.4.4 Benchmark Against Low- and Mid-Level Movie Features

The results of the correlation analysis with the low- and mid-level movie features are shown in Figure 5.7, Figure 5.8, Figure 5.9 and Figure 5.10 respectively.

In terms of luminance, our observations reveal contrasting patterns between Sherlock and Iteration in regions of the visual cortex. While Sherlock demonstrates a positive correlation, indicating a relationship between luminance and brain activity, Iteration exhibits a negative correlation in the same regions. Conversely, when examining changes in luminance (i.e., the derivative of the luminance time course), both datasets display a positive correlation within the visual cortex. These findings can be interpreted as a consequence of the

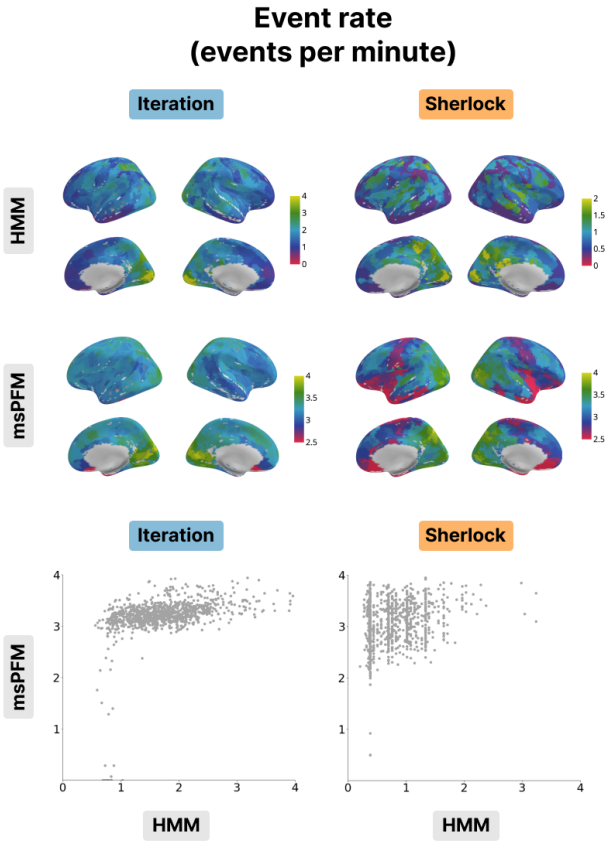

Figure 5.5: Event rate for the Iteration (left) and Sherlock (right) datasets. The event rate maps obtained with the HMM (top) and msPFM algorithms (bottom) are compared. The bottom row shows the differences in values between msPFM and the HMM approach.

luminance adaptation in Iteration. The movie predominantly features very bright scenes, occasionally transitioning to dark moments as shown in the blue time courses at the bottom left of Figure 5.7. In contrast, Sherlock contains more diverse scenes in multiple locations, thereby exhibiting a greater range of luminance with both the luminance and its derivative being very similar (see red time courses in the bottom left of Figure 5.7).

The audio envelope analysis reveals a noteworthy positive correlation within regions of the auditory cortex for both datasets, with a more pronounced correlation observed in Iteration. Examination of audio envelope variations, i.e., the derivate of the audio envelope, further demonstrates a robust positive correlation within the auditory cortex for both datasets. In this case, both the audio envelope and its derivate showed a very similar pattern (see red and blue

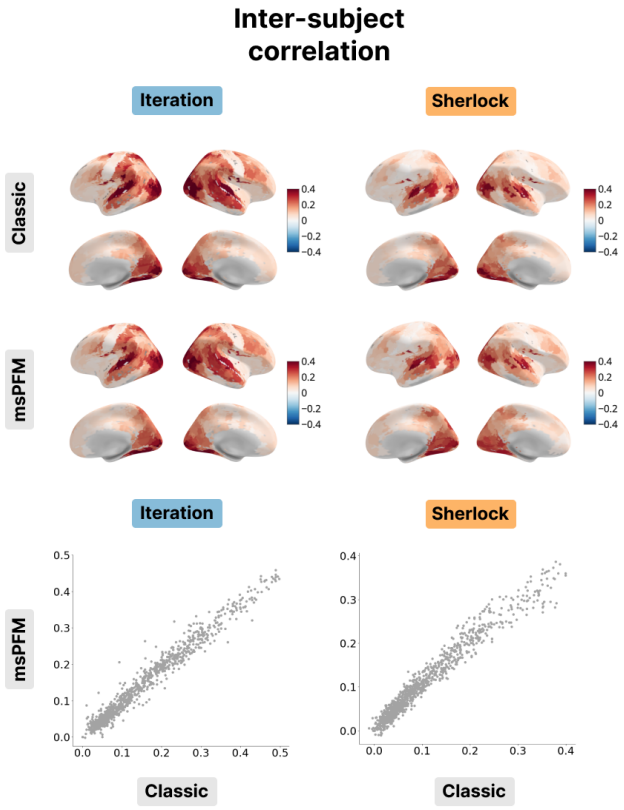

Figure 5.6: Inter-subject correlation for the Iteration (left) and Sherlock (right) datasets. The inter-subject correlation obtained with msPFM (bottom) is compared with the inter-subject correlation obtained with the classic approach (top). The bottom row shows the differences in values between msPFM and the classical ISC approach.

time courses in the bottom right of Figure 5.8). This underscores the ability of msPFM to accurately capture the intricate neural dynamics associated with the auditory component of the movies.

Next, mid-level movie features were examined, including the presence of speech as well as hands and faces onscreen. Significantly, it can be observed that msPFM adeptly captures the intricate neural dynamics linked to speech in both datasets, revealing large correlations in expected areas of the superior temporal gyrus, Heschl's gyrus, middle temporal and inferior frontal gyrus, all of them bilaterally. These maps demonstrate the adaptability of msPFM to capture varying amounts of speech present in the respective movies. This adaptability is particularly pronounced in the context of Iteration, characterized by a scarcity of speech compared to Sherlock—a movie abundant in speech- as

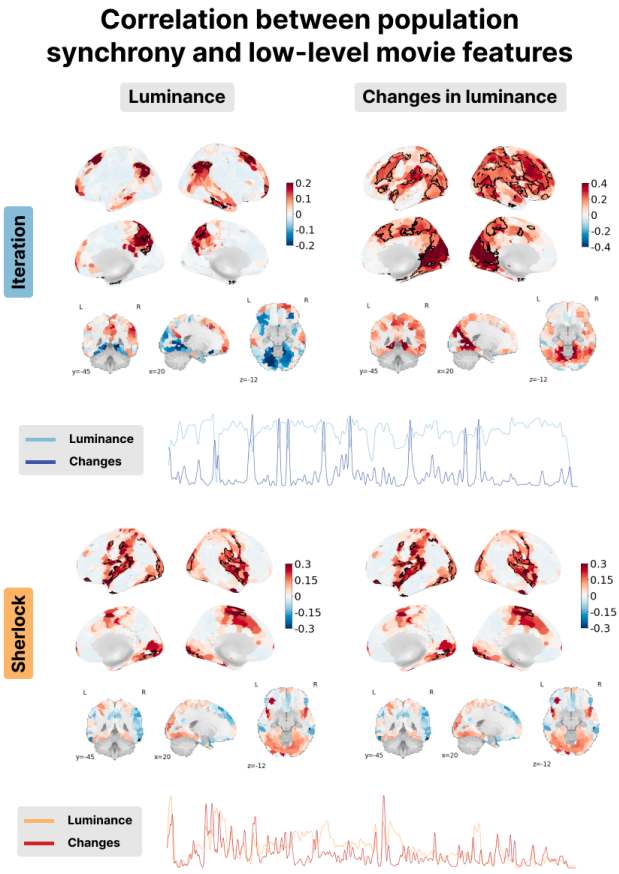

Figure 5.7: Correlation between the PopSync+ and low-level features of the movie with luminance on the left and its derivative on the right. Significant regions are highlighted by a black contour, while non-significant regions are displayed with increasing transparency as the correlation values are further from the significance threshold. The correlation maps are shown for the Iteration (top) and Sherlock (bottom) datasets. The time courses below depict the luminance (lighter color) and its derivate (darker color) for Iteration (blue) and Sherlock (red).

shown in the time courses in the bottom left of Figure 5.9. Notwithstanding this discrepancy, msPFM reliably captures the neural dynamics associated with speech in language-related areas, reaffirming its efficacy.

Furthermore, positive correlation between the msPFM-estimated activity-inducing signal and the presence of hands in the movies can also be observed in bilateral sensorimotor regions around the central sulcus. In the context of Iteration, the correlation was observed prominently within regions of the

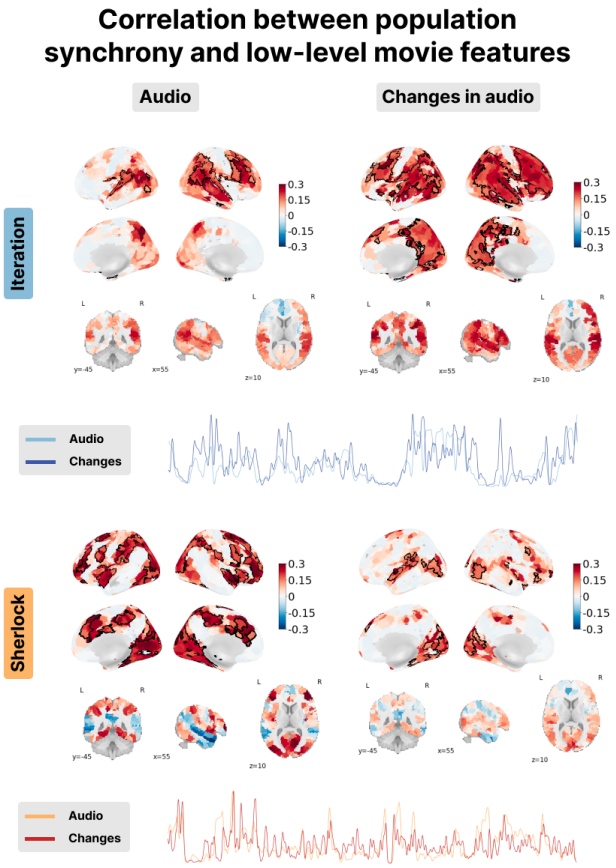

Figure 5.8: Correlation between the PopSync+ and low-level features of the movie with luminance and its derivative on the left, and the audio envelope and its derivative on the right. Significant regions are highlighted by a black contour, while non-significant regions are displayed with increasing transparency as the correlation values are further from the significance threshold. The correlation maps are shown for the Iteration (top) and Sherlock (bottom) datasets. The time courses of the audio envelope (lighter color) and its derivate (darker color) are shown in blue for Iteration and in red for Sherlock.

somatosensory cortex (i.e. postcentral gyrus) and supplementary motor areas, while Sherlock exhibited correlation within more anterior regions in the motor cortex (i.e. precentral gyri). These findings could potentially be explained by the contrasting composition of Iteration and Sherlock. While Iteration predominantly incorporates hands in relation to other objects or the main character's actions (see the high frequency in the blue time course at the bottom left of Figure 5.9), it features fewer hands-only shots. On the other hand, Sherlock exhibits a higher

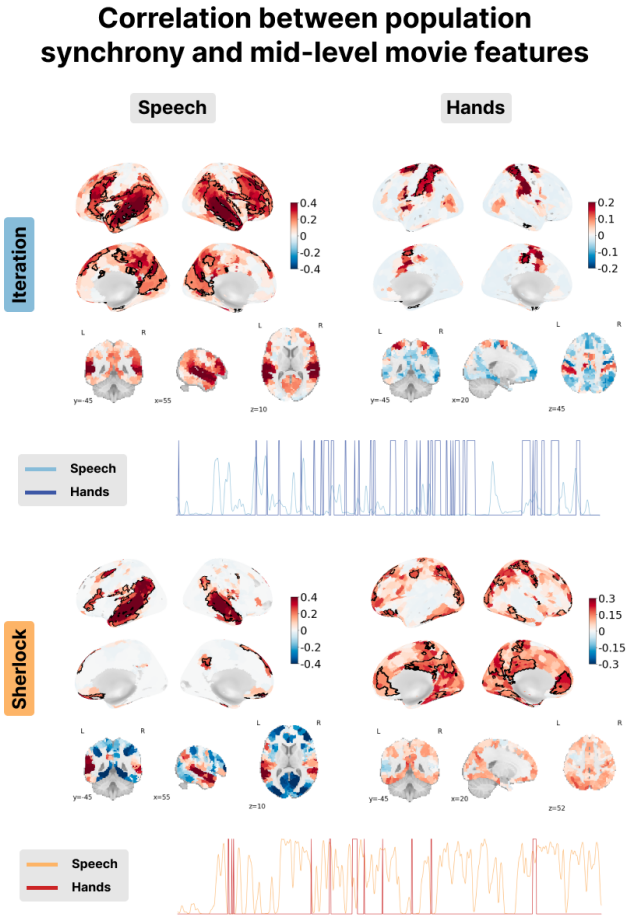

Figure 5.9: Correlation between the PopSync+ and mid-level features of the movie with the presence of speech on the left and hands on the right. Significant regions are highlighted by a black contour, while non-significant regions are displayed with increasing transparency as the correlation values are further from the significance threshold. The correlation maps are shown for the Iteration (top) and Sherlock (bottom) datasets. The time courses below depict the presence of speech (lighter color) and hands (darker color) for Iteration in blue and Sherlock in red.

frequency of hands-only shots (though few across the entire movie as shown in red in the bottom left time course in Figure 5.9), which may account for the aforementioned findings.

Finally, as shown in Figure 5.10, correlation between the PopSync+ time-series and the presence of faces was found in the Iteration movie in the left

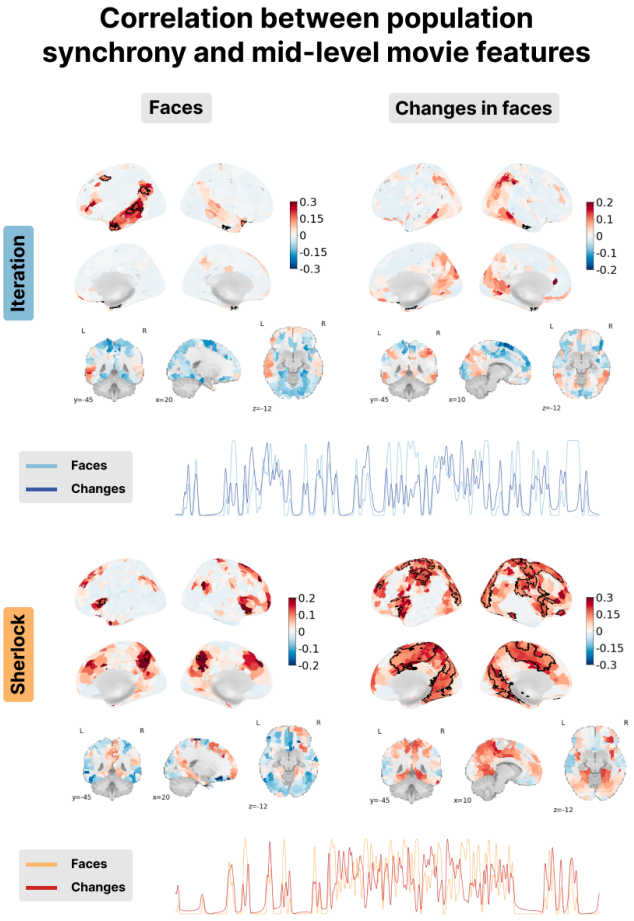

Figure 5.10: Correlation between the PopSync+ and mid-level features of the movie with the presence of faces on the left and its derivative on the right. Significant regions are highlighted by a black contour, while non-significant regions are displayed with increasing transparency as the correlation values are further from the significance threshold. The correlation maps are shown for the Iteration (top) and Sherlock (bottom) datasets. The time courses of the presence of faces (lighter color) and its derivate (darker color) are shown in blue for Iteration and red for Sherlock.

middle temporal gyrus, a region associated with facial familiarity (Zhen et al., 2013), and the posterior superior temporal sulcus, a region associated with the processing of gaze and expression (Baseler et al., 2012). This finding aligns with the movie's nature, as Iteration predominantly features a single character and contains several headshots where viewers have their eyes locked on the character's

eyes and can see her facial expression. No significant or substantial correlation between the PopSync+ measure and the presence of faces was found in the movie in the Sherlock dataset. However, when focusing on changes in faces specifically within the Sherlock dataset, a positive correlation was observed in the fusiform gyrus, a region associated with face recognition (Kanwisher & Yovel, 2006; Kanwisher et al., 1997; Rossion et al., 2003) Interestingly, this correlation with changes in faces was not as strong when examining changes in faces for Iteration. These contrasting findings may be attributed to the distinctive characteristics of the movies. In Sherlock, which showcases multiple characters and their facial expressions during dialogue, the adaptation or identification of faces may play a role (see long periods of changing faces in the red time courses at the bottom of Figure 5.10). On the contrary, Iteration primarily focuses on a single character, suggesting an fMRI adaptation effect in fusiform face areas, i.e. a lower response to an identically repeated face than to new faces (Avidan & Behrmann, 2010; Eger et al., 2005; Gauthier & Nelson, 2001; Pourtois et al., 2005; Rotshtein et al., 2004; Yovel & Kanwisher, 2004).

Overall, these analyses showcase the potential of msPFM in capturing the activity linked to lower- and mid-level stimulus features within anticipated brain regions and timeframes, even without prior knowledge of event timings and regardless of the stimulus itself. Furthermore, these findings highlight the ability of msPFM to identify additional stimulus-specific regions associated with the intricate and distinctive content inherent to each stimulus.

5.5 Discussion

The potential of multi-subject Paradigm Free Mapping (msPFM) to detect and elucidate moment-to-moment spatio-temporal neural activity patterns evoked by naturalistic stimuli has been showcased using simulations and two experimental datasets. While previous investigations have employed alternative approaches, such as hidden Markov models (HMM) (Baldassano et al., 2017), greedy state boundary search (GSBS) (Geerligs et al., 2021), inter-subject correlation (Nastase et al., 2019) or dynamic functional connectivity (Di & Biswal, 2020), these methods operate on a limited number of regions of interest (ROIs) and condense the temporal dynamics of the BOLD signal into states or connectivity patterns that encompass a time period exceeding the acquired data's temporal resolution. Consequently, the results shown here highlight the unparalleled potential of msPFM to operate at the utmost temporal and spatial precision achievable.

Importantly, msPFM was validated by comparing it against other common algorithms used to assess naturalistic fMRI data, such as Hidden Markov Models (HMM) and intersubject correlation (ISC). First, the event-rate analysis shown in Figure 5.5 revealed a distinct cortical gradient, characterized by faster event rates in posterior sensory regions and slower rates in anterior higher-order regions. These results align with previous findings on neural event boundaries across subjects and the segmentation of events using HMM (Baldassano et al., 2017; Sava-Segal et al., 2022), which could support the idea that the brain's cortical

organization is arranged hierarchically, with sensory inputs passing through unimodal areas and being abstracted into broader conceptual and cognitive representations in transmodal areas (Bernhardt et al., 2022; Margulies et al., 2016; Samara et al., 2023). However, msPFM offers several advantages over the HMM-based approach. First, the HMM method requires specifying a desired number of events (k) for each region. While there are data-driven ways to estimate an appropriate value for k at the group level, these require additional computational time and power, and are not straightforward to implement for individual subjects. In contrast, msPFM does not require pre-specifying any value of k, and instead recovers individual neuronal-related events at both the group and individual level. Furthermore, the HMM model assumes that each "state" exhibits a unique neural activity signature that changes at state boundaries, which may not hold true universally, as certain brain regions could potentially revert to previous or new states inside this perdiod. In contrast, msPFM applies no restriction on the number and duration of the neural activity patterns that can be detected, thereby providing a more comprehensive representation of the brain's spatio-temporal dynamics. The crucial distinction is that the HMM approach identifies events or states that persist for a specific duration, whereas the msPFM technique retrieves discrete instances of neuronal-related BOLD activity, which can be directly associated with individual moments in the stimulus. Another important distinction is the computational cost needed to analyse data with each of the two techniques. In fact, while finding the appropiate number of events with HMM can take >24h at the single-subject level with 1000 ROIs using a high performance computing server with up to 512GB of RAM available, finding estimates of the activity-inducing signal and counting the number of events with msPFM only took between 5 and 10 minutes.

Furthermore, techniques like HMM and inter-subject correlation are not specifically designed to detect activity evoked by the stimulus in different brain regions, but rather to identify regions that encode stimulus-related information consistently across multiple individuals (Nastase et al., 2019). In contrast, msPFM can detect moment-to-moment neuronal resonses not only shared across subject, but also distinguish patterns occurring in individual subjects. By correlating the recovered population synchrony with features of the movies, ?? and ?? show that msPFM is able to detect neuronal activity driven by both low- and mid-level features of the stimulus in both primary sensory and higher-order brain regions. Particularly, fMRI adaptation to the stimulus (Grill-Spector & Malach, 2001) may be present in some features such as luminance, with the visual cortex mostly showing activity in response to changes in the luminance in Iteration probably due to the increased contrast associated with these changes (Goodyear & Menon, 1998); or the presence of faces, with Iteration having a single character and no reliable face-related activity in the fusiform gyrus as opposed to Sherlock and the many characters in the clip, where the fusiform gyrus showed correlation with changes in faces (Kanwisher & Yovel, 2006; Kanwisher et al., 1997; Rossion et al., 2003). The case of Sherlock is particularly challenging because faces are always closely followed by speech, which is a stronger effect and can hide the response to the presence of faces (de la Vega et al., 2022).

While this study introduces a novel approach for analyzing naturalistic fMRI data, there is potential for further enhancements. Adopting a data-driven strategy like stability selection (Meinshausen & Bühlmann, 2010; Uruñuela et al., 2022) could eliminate the need for manually selecting the regularization parameter. By adopting such an approach, not only would the method's robustness be enhanced, but it would also enable estimation of the probability of neuronal-related events at each TR, subject, and voxel or ROI. Additionally, the utilization of optimal transport methods could further refine the spatial differentiation of shared and individualized patterns by improving its robustness against anatomical (spatial) and hemodynamic (temporal) variability between subjects (Janati et al., 2020). Nevertheless, the current methodology adequately demonstrates the potential of msPFM-alike approaches in investigating the brain's spatio-temporal dynamics during naturalistic stimuli. For instance, msPFM offers a valuable approach to unraveling the precise neural mechanisms involved in event segmentation. It can be employed independently or in conjunction with other techniques such as HMM or GSBS to gain an understanding of the timescale for memory encoding, as well as the continuous updating or encoding of memory fragments following the completion of an event (Baldassano et al., 2017; Silva et al., 2019). Moreover, novel research utilizing msPFM could be aimed at comprehending how the human brain comprehends these complex multiplexed signals and understanding the specific features of the stimulus that elicit responses in the human brain. Finally, the proposed method has the potential to elucidate the connection between individual differences in these responses and subsequent memory formation or appraisal of the stimulus.

5.6 Code and Data Availability

The code and materials used to generate the figures in this work can be found in the following GitHub repository: https://github.com/eurunuela/msPFM_paper.

The msPFM Python package is available in the following GitHub repository: https://github.com/ParadigmFreeMapping/msPFM

Chapter

6

Sparse and Low-Rank Paradigm Free Mapping

This chapter was published as Uruñuela, E., Moia, S., & Caballero-Gaudes, C. (2021). A low rank and sparse paradigm free mapping algorithm for deconvolution of fMRI data. 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI). DOI: https://doi.org/10.1109/ISBI48211.2021.9433821.

Current deconvolution algorithms for functional magnetic resonance imaging (fMRI) data are hindered by widespread signal changes arising from motion or physiological processes (e.g. deep breaths) that can be interpreted incorrectly as neuronal-related hemodynamic events. This work proposes a novel deconvolution approach that simultaneously estimates global signal fluctuations and neuronal-related activity with no prior information about the timings of the blood oxygenation level-dependent (BOLD) events by means of a sparse and low rank decomposition algorithm. The performance of the proposed method is evaluated on simulated and experimental fMRI data, and compared with state-of-the-art sparsity-based deconvolution approaches and with a conventional analysis that is aware of the temporal model of the neuronal-related activity. This chapter demonstrates that the novel Sparse and Low Rank Paradigm Free Mapping (SPLORA-PFM) can estimate global signal fluctuations related to motion in the task, while estimating the neuronal-related activity with high fidelity. The open-source Python package for SPLORA is available at https://github.com/ParadigmFreeMapping/splora

6.1 Introduction

As noted in previous chapters of this thesis, hemodynamic deconvolution algorithms of functional magnetic resonance imaging (fMRI) data aim to estimate blood oxygenation level-dependent (BOLD) events with no prior knowledge of their timing. These algorithms can be specially useful when the information about the timing of the neuronal activity that drives the BOLD events is unknown, inaccurate or insufficient (e.g., resting-state, naturalistic paradigms, clinical conditions). However, the performance of existing deconvolution approaches can be hampered considerably in presence of global, widespread signal changes due to head jerks, hardware artefacts or prominent non-neuronal physiological events (e.g., deep breaths) (Power et al., 2017). Signal artefacts due to head motion and hardware malfunction can be reduced by means of denoising algorithms, such as ICA-AROMA (Pruim, Mennes, van Rooij, et al., 2015) or ME-ICA (Kundu et al., 2012), or can be compensated with a multi-echo Paradigm Free Mapping formulation (Caballero-Gaudes et al., 2019). However, global physiological events are more difficult to compensate during data preprocessing (Power et al., 2018) and can be misinterpreted as neuronally related since their temporal signature can closely resemble the hemodynamic response function (HRF) assumed in the deconvolution model to describe neurovascular coupling.

This chapter proposes a new Paradigm Free Mapping algorithm for spatio-temporal deconvolution of fMRI data that is capable of simultaneously estimating global signal fluctuations and neuronal-related activity based on a sparse and low-rank decomposition approach. The proposed algorithm extends the formulation of the multivariate-sparse Paradigm Free Mapping (Mv-SPFM) introduced in Chapter 4 by using a regularized estimator consisting of the same structured sparsity promoting $\ell_{2,1}+\ell_1$ norm but adding a low-rank-promoting nuclear-norm (Otazo et al., 2015).

6.2 Sparse and Low-Rank Paradigm Free Mapping

Let us consider that the whole-brain fMRI data $\mathbf{Y} \in \mathbb{R}^{N \times V}$ where N is the number of volumes and V is the number of voxels of the acquisition can be decomposed into three terms, i.e.

$$\mathbf{Y} = \mathbf{HS} + \mathbf{L} + \mathbf{N}, \tag{6.1}$$

where the neuronal-related component $\mathbf{HS}$ is the convolution of voxel-specific neuronal-related signals $\mathbf{S}$ with the Toeplitz matrix $\mathbf{H} \in \mathbb{R}^{N \times N}$ with shifted HRFs in its columns (i.e. similar to the formulation used for multivariate Paradigm Free Mapping Uruñuela et al., 2022), the global fluctuations can be captured as the sum of P spatially widespread (i.e. global) low-rank components $\mathbf{L} = \sum_{p=1}^{P} \mathbf{v}_p \mathbf{a}_p^T$ where $\mathbf{v}_p \in \mathbb{R}^{N \times 1}$ and $\mathbf{a}_p \in \mathbb{R}^{V \times 1}$ denote their corresponding spatial and temporal signatures, and $\mathbf{N}$ represents additional white Gaussian noise.

The following multivariate regularized least-squares problem is proposed to

Algorithm 1 SPLORA-PFM algorithm using MFISTA-VA

input: $\mathbf{Y}, \mathbf{H}$
initialize: $\mathbf{L_0}, \mathbf{S_0}, \mathbf{Y_{S,0}}, \mathbf{Y_{L,0}}, \mathbf{Y_{A,0}} = 0,\ c = \|\mathbf{H}\|_F^2$
while not converged do
    $\mathbf{Z_S} = \mathbf{Y_{S,k}} + (\mathbf{1}/\mathbf{c}) * (\mathbf{Y} - \mathbf{Y_{A,k}})$
    $\mathbf{Z_L} = \mathbf{Y_{L,k}} + (\mathbf{1}/\mathbf{c}) * (\mathbf{Y} - \mathbf{Y_{A,k}})$
    # $\mathbf{L}$: singular-value soft thresholding (SVT)
    $\mathbf{L_k} = \mathrm{SVT}_{\lambda_L}(\mathbf{Z_S})$
    # $\mathbf{S}$: proximity operator for the $\ell_{2,1}$+ℓ_1 norm
    $\mathbf{S_k} = \mathrm{prox}_{\mathbf{D_S}}(\mathbf{Z_L})$
    # Update $\mathbf{A}$
    $\mathbf{Z_A} = (\mathbf{L_k} - \mathbf{Y_{L,k}}) + \mathbf{H}(\mathbf{S_k} - \mathbf{Y_{S,k}})$
    $\mathbf{A_k} = \mathbf{Y_{A,k}} + \mathbf{Z_A}$
    # Calculate MFISTA step size: $t_k = \frac{1+\sqrt{1+4*t_{k-1}^2}}{2}$
    # Calculate η_k as in (Zibetti et al., 2018)
    $\mathbf{Y_{S,k+1}} = \mathbf{S_k} + \frac{t_k - 1}{t_{k+1}}(\mathbf{S_k} - \mathbf{S_{k-1}}) + \frac{t_k}{t_{k+1}}(\mathbf{Z_S} - \mathbf{S_k}) + \frac{t_k}{t_{k+1}}(\eta_k - 1)(\mathbf{Z_S} - \mathbf{Y_{S,k}})$
    $\mathbf{Y_{L,k+1}} = \mathbf{L_k} + \frac{t_k - 1}{t_{k+1}}(\mathbf{L_k} - \mathbf{L_{k-1}}) + \frac{t_k}{t_{k+1}}(\mathbf{Z_L} - \mathbf{L_k}) + \frac{t_k}{t_{k+1}}(\eta_k - 1)(\mathbf{Z_L} - \mathbf{Y_{L,k}})$
    $\mathbf{Y_{A,k+1}} = \mathbf{A_k} + \frac{t_k - 1}{t_{k+1}}(\mathbf{A_k} - \mathbf{A_{k-1}}) + \frac{t_k}{t_{k+1}}(\mathbf{Z_A} - \mathbf{A_k}) + \frac{t_k}{t_{k+1}}(\eta_k - 1)(\mathbf{Z_A} - \mathbf{Y_{A,k}})$
end while
output: $\mathbf{L_k}, \mathbf{S_k}$

estimate both the neuronal-related signals and the global components:

$$\hat{\mathbf{L}}, \hat{\mathbf{S}} = \arg\min_{\mathbf{L},\mathbf{S}} \|\mathbf{Y} - \mathbf{H}\mathbf{S} - \mathbf{L}\|_F^2 + \lambda_L \|\mathbf{L}\|_* + (1-\rho)\|\mathbf{D}_S\mathbf{S}\|_{2,1} + \rho\|\mathbf{D}_S\mathbf{S}\|_1, \tag{6.2}$$

where $\|\cdot\|_F$ denotes the Frobenious norm, the $\ell_{2,1}$+ℓ_1-norm term enforces temporal sparsity and spatial structure on the estimate of the neuronal-related activity and ρ controls the tradeoff between both terms (Gramfort et al., 2011; Uruñuela et al., 2022) and $\mathbf{D}_s = \mathrm{diag}\left(\lambda_{S_1}, \ldots, \lambda_{S_V}\right)$ is a diagonal matrix with voxel-specific non-negative regularization parameters that balances the sparsity of $\mathbf{S}$ and data fidelity for each voxel. In addition, the nuclear-norm $\|\cdot\|_*$ encourages the estimation of low-rank components where the non-negative regularization parameter λ_L controls the number of low-rank components.

Here, ρ is empirically set to 0.8 to enforce structure in the spatial domain and maintain the sparsity of the estimates. For each voxel, λ_{S_i} is set equal to the median absolute deviation estimate of the noise standard deviation from the fine-scale wavelet coefficients of the voxel time series (Daubechies, order 3). After the singular value decomposition (SVD) of the data, λ_L is set to select P low-rank components corresponding to the largest eigenvalues showing a difference of at least 10% with respect to the next eigenvalue.

The optimization problem in Equation (6.2) is solved via monotone FISTA with variable acceleration (MFISTA-VA) (Zibetti et al., 2018) as shown in Algorithm 1.

6.3 Methods

6.3.1 Simulated Data

1000 voxels were simulated including two groups of 50 voxels with a known BOLD signal, whereas the remaining voxels did not contain any BOLD signal. For each voxel, different signal sources representing motion-related, thermal and physiological noise were added following (Gaudes et al., 2013), as well as two global low-rank components (see Figure 6.1 A) with a random voxelwise amplitude simulating widespread signal changes due to two deep breaths (Power et al., 2018) and large amplitude spikes mimicking spin-history artefacts due to head jerks, respectively.

The performance of the proposed SPLORA-PFM algorithm was asssesed on different signal to noise ratio (SNR) settings and with different ratios of voxels with BOLD signals to total number of voxels (denoted as BOLD/total voxels ratio). Four different algorithms based on Equation (6.2) were evaluated depending on the regularization parameters:

1. SPFM with no low-rank estimation and no spatial regularization (SPFM, $\rho = 1,\ \lambda_L = 0$)
2. MV-SPFM with no low-rank estimation (MV-SPFM, $\rho = 0.8,\ \lambda_L = 0$)
3. SPLORA-PFM algorithm with only the L1-norm (LR+SPFM, $\rho = 1$)
4. SPLORA-PFM algorithm ($\rho = 0.8$)

These were benchmarked against the original univariate SPFM algorithm with regularization parameter selected according to the Bayesian Information Criterion (SPFM-BIC) (Gaudes et al., 2013). Note that both the SPFM algorithm with $\rho = 1$ and $\lambda_L = 0$ and the original SPFM-BIC algorithm operate voxelwise, except the regularization parameters are chosen differently.

True and false positive and negative values were calculated for the ROC values comparing the estimated activity-inducing signal with the simulated (binary) activity-inducing signal (ground truth) as follows: a TR was deemed as a true positive (TP) when the estimated and simulated values were both non-zero; a TR was treated as true negative (TN) when both the estimated and simulated values were zero; false positives (FP) were given to those TRs with a non-zero estimated value when the simulated value was zero; and false negatives (FN) were considered when the estimated value was zero but the simulated value was non-zero. Sensitivity and specificity values were calculated as $\mathrm{TP}/(\mathrm{TP}+\mathrm{FN})$ and $\mathrm{TN}/(\mathrm{TN}+\mathrm{FP})$ respectively.

6.3.2 Experimental Data

Nine healthy subjects were scanned in a 3T Siemens Prisma MR scanner in ten sessions at the same hour and day of the week. T2*-weighted multi-echo fMRI data was collected with a multiband (MB) multiecho gradient echo planar

imaging sequence (340 scans, 52 slices, Partial-Fourier=6/8, voxel size=2.4x2.4x3 mm^3, TR=1.5 s, TEs=10.6/28.69/46.78/64.87/82.96 ms, flip angle=70^o, MB factor=4, GRAPPA=2). During the fMRI acquisition, subjects performed a motor task consisting of five different movements (left-hand finger tapping, right-hand finger tapping, moving the left toes, moving the right toes and moving the tongue). These conditions were randomly intermixed every 16 seconds, and were only repeated once the entire set of conditions were presented. For this work, only the first two sessions were selected to evaluate the algorithm.

Data preprocessing was carried out with AFNI (Cox, 1996) including volume realignment, optimally combining the echo time datasets with tedana (DuPre et al., 2021), detrending of up to 5^{th}-order Legendre polynomials, spatial smoothing with a Gaussian kernel of 3 mm Full Width Half Maximum, and normalization to signal percentage change. Based on the simulation results, preprocessed data was analyzed with the novel SPLORA-PFM algorithm with $\rho = 0.8$, and λ_L and $\mathbf{D}_S$ were selected as described in Section 6.2.

6.4 Results and Discussion

6.4.1 Simulated Data

Figure 6.1B depicts the receiver operating characteristic (ROC) curves with the sensitivity and specificity rates for the estimation of the neuronal-related signal $\hat{\mathbf{S}}$ for each simulation scenario. Regardless of the simulated SNR and the BOLD/total voxels ratios, the ROC values demonstrate the proposed SPLORA-PFM algorithm achieves higher specificity and sensitivity than the original SPFM method, except for the highest BOLD/total number of voxels ratio and highest SNR where the multivariate nature of the model prevents the algorithm from fitting accurately each voxel. As expected, all variations of the proposed algorithm exhibit lower sensitivity as the SNR is reduced while maintaining the level of specificity. In addition, Figure 6.1C plots the error of the low-rank component estimate obtained with the SPLORA-PFM algorithm for $\rho = 0.8$, showing that its estimate improves with a lower BOLD/total voxels ratio.

6.4.2 Experimental Data

Figure 6.2A-F depict the results of the SPLORA-PFM algorithm in a representative dataset. For this subject and session, the proposed method estimated $P = 3$ global low-rank components whose time series $\mathbf{a}_p$ and spatial maps $\mathbf{v}_p$ are shown in Figures 2C and 2F, respectively. Figure 6.2A shows the Euclidean norm of the motion displacements (E-norm), DVARS (the spatial root mean square of the data; Power et al., 2012; Smyser et al., 2011) and the average global signal (GS) time series, whereas Figure 6.2B displays the grayplots of the preprocessed data (RAW), estimated low-rank component and estimated neuronal-related component in gray matter (GM) and white matter (WM) voxels. The first low-rank component captures signal fluctuations related to head movements and susceptibility artefacts during the 'moving the tongue'

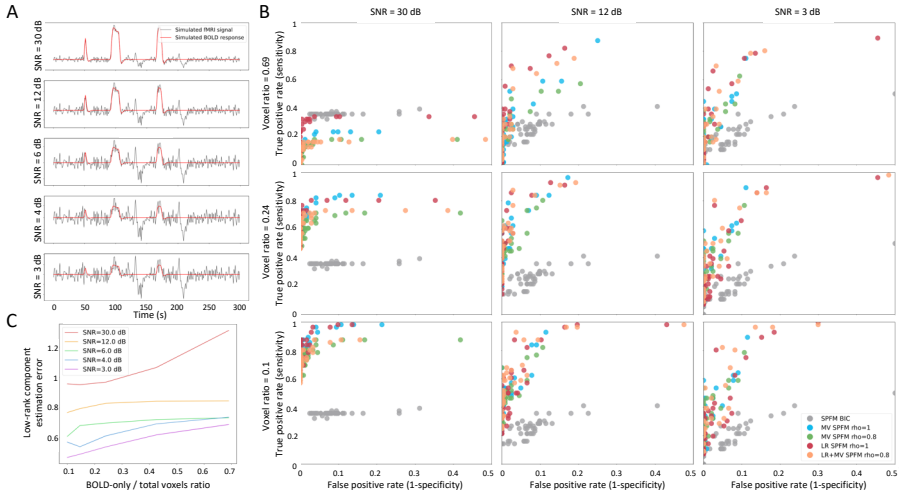

Figure 6.1: Simulation results. A) An example of the simulated signals for the different SNR conditions; B) ROC values for the estimation of the neuronal-related signal with: SPFM using BIC (SPFM BIC), SPFM with no low-rank estimation and no spatial regularization (SPFM, $\rho = 1$), MV-SPFM with no low-rank estimation (MV-SPFM, $\rho = 0.8$), the SPLORA-PFM algorithm with only the L1-norm (LR+SPFM, $\rho = 1$), and the SPLORA-PFM algorithm ($\rho = 0.8$). C) Estimation error of the low-rank components for different ratios of BOLD/total number of voxels.

condition, suggesting that the subject moved the head while performing the tongue movement task. The second low-rank component has a time series that closely follows the global signal and its spatial map actually delineates major arteries and draining veins, whereas the third component is clearly related to widespread physiological fluctuations. Among participants, the number of estimated low-rank components ranged between 1 and 5.

Furthermore, Figure 6.2D and Figure 6.2E illustrate the time series of the estimated neuronal-related signal for a representative voxel (see cross in the first map) and the maps for several individual events of the tongue movements and right hand finger tapping conditions, respectively. The SPLORA-PFM maps reveal clusters of activity in similar regions to those inferred with a traditional general linear model (GLM) analysis, which is aware of the timings of the events. The single-trial GLM activation maps are thresholded based on their t-statistic at a significance threshold of $p < 0.001$. Notably, the SPLORA-PFM maps still depict the tongue areas of the motor cortex bilaterally even though the timing of the first low-rank component also closely followed the tongue condition.

Finally, Figure 6.3 depicts the ROC values of the MV-SPFM and SPLORA-PFM, both using $\rho = 0.8$, and the original SPFM algorithm using the GLM maps of each event thresholded at a $p = 0.001$ as the ground-truth. The ROC

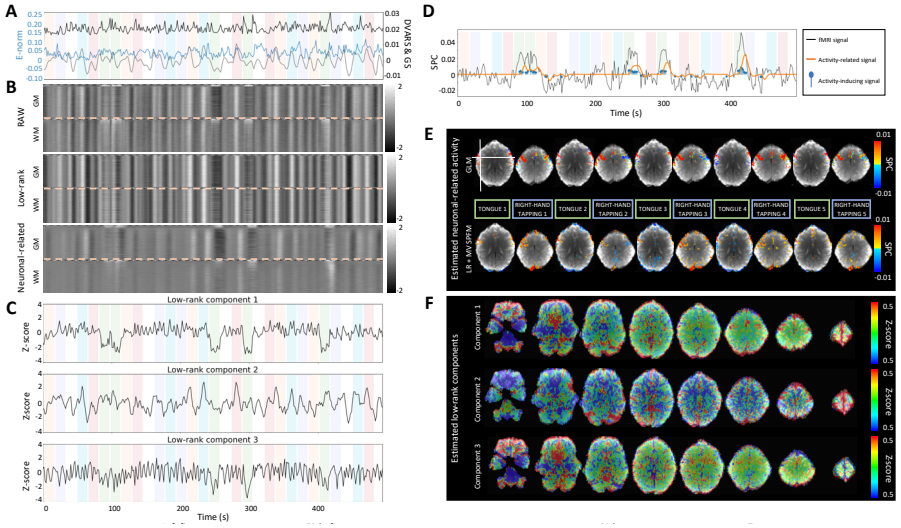

Figure 6.2: A) Euclidean norm of the motion displacements (E-norm) (blue), DVARS (black) and global (gray) signals of the fMRI data; B) Grayplots of gray matter (GM) and white matter (WM) of the preprocessed fMRI data, the estimated low-rank and neuronal-related components; C) Time series and F) maps of the estimated low-rank components; D) Time series and E) maps of representative single-trial neuronal-related (motor) activity. The color bands in the plots with the time series illustrate the timing of the different conditions.

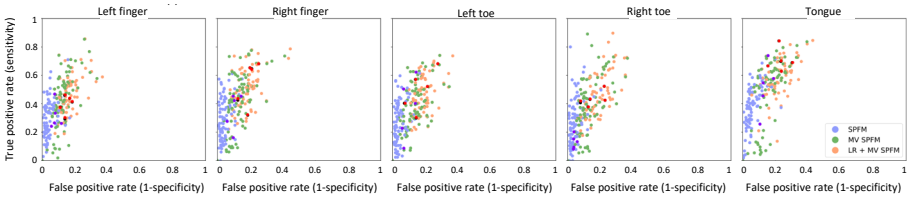

Figure 6.3: ROC values of the five conditions for the three algorithms tested: SPFM, MV-SPFM and SPLORA-PFM (red, dark-purple and dark-green dots correspond to the subject in Figure 6.2).

curves of the five motor task conditions show that both the MV-SPFM and the SPLORA-PFM approaches provide higher sensitivity at the cost of a reduced specificity, and that the higher complexity involved in estimating the low-rank component does not diminish the accuracy in deconvolving the neuronal-related component of the signal.

This work introduced a novel formulation for the deconvolution of BOLD fMRI data using a low rank and sparse algorithm that captures global fluctuations

due to motion artefacts or physiological signals that typically reduce the accuracy of neuronal related estimates of currently used algorithms. The formulation described in this chapter was presented for single-echo acquisition and the spike model. However, it can be adapted straightforwardly for multi-echo (Caballero-Gaudes et al., 2019) or the block model (Cherkaoui et al., 2019; Karahanoğlu et al., 2013; Uruñuela et al., 2023) using the multi-echo and innovation signals described in Chapter 4 and Chapter 2 respectively. Likewise, the selection of the regularization parameters $\mathbf{D}_s$, λ_L and ρ was done empirically here and could be further optimized using robust approaches such as stability selection (Meinshausen & Bühlmann, 2010; Uruñuela et al., 2020) as described in Chapter 3 and Chapter 4. Finally, future work will be directed towards evaluating the performance of SPLORA-PFM on resting-state data. Specifically, these evaluations will assess its effectiveness in mitigating the influence of commonly observed global fluctuations on activity-inducing signal estimates, since it is expected that SPLORA-PFM will exhibit superior performance in this regard compared to its performance on task fMRI data.

Chapter

7

Conclusion and Future Directions

7.1 Summary and Final Discussion

As mentioned in the introduction (Chapter 1), conventional single trial fMRI analysis relies heavily on preexisting knowledge of event timing (Buckner et al., 1996; Menon et al., 1998; Richter et al., 1997). However, in certain situations such as clinical settings, resting-state, or naturalistic paradigms, it can be challenging or even impossible to formulate a temporal model of activations. This thesis has introduced new fMRI analysis techniques that allow for the study of true single trial BOLD responses without any prior information about the event timings. These methods greatly improve upon the Paradigm Free Mapping approach (Gaudes et al., 2013), and extend its capabilities to include multivariate and multi-subject settings.

In Chapter 2, the underlying motivations and principles of Paradigm Free Mapping (PFM) and hemodynamic deconvolution are discussed. Additionally, the chapter also delved into existing fMRI analysis methods that aim to conduct single trial experiments without the need for precise or null specifications of the experimental paradigm. Notably, Total Activation (TA) was highlighted as a leading deconvolution technique that follows an analysis formulation, in contrast to the synthesis formulation of PFM. The chapter also introduced the block model for PFM, which allows for the estimation of the innovation signal –the derivative of the activity-inducing signal. This approach is especially useful for block-design experiments and utilizes the sparsity constraint of the LASSO more effectively. After thorough comparison, it was found that the two methods are essentially equivalent, with PFM emerging as the preferred option due to its simplistic and adaptable formulation. This paved the way for the development of various PFM techniques in this thesis, including stability-based PFM (Chapter 3), multivariate PFM (Chapter 4), multi-subject PFM (Chapter 5), and sparse and low-rank PFM (Chapter 6).

A crucial aspect for achieving satisfactory operation of the PFM techniques is the accurate selection of the regularization parameters, as the deconvolution relies on regularized estimators. In the original PFM techniques (Gaudes et al., 2013), the choice of the regularization parameter λ was determined using Akaike and Bayesian information criteria independently for each voxel. Chapter 3 introduced an alternative approach based on stability selection, which eliminates the need for selecting this parameter altogether. Moreover, this new procedure offers an additional metric, defined as the area under the curve of the stability path, that represents the likelihood of the detected events being true at the finest spatial and temporal scales.

Chapter 4 presented a novel approach to the PFM methodology, introducing a multivariate formulation that organizes voxel or ROI time series into a matrix and incorporates the stability selection procedure. By combining these two techniques with the use of the $\ell_1 + \ell_{2,1}$ mixed norm group sparsity regularization –which incorporates spatial information and interactions into the formulation– the chapter showcased the enhanced performance of the multivariate PFM with stability selection in comparison to the original univariate formulation of PFM. Notably, this method yielded results that were more concordant with single-trial GLM findings. Additionally, the chapter demonstrated the robustness of the approach across both single-echo and multi-echo fMRI data. This was achieved through the utilization of the area under the curve (AUC) measure, which facilitated the application of various thresholding techniques adaptable to the noise level inherent in the data.

Chapter 5 introduced a new application of the multivariate PFM formulation: simultaneous deconvolution of multiple subjects performing a naturalistic paradigm. In this case, instead of concatenating voxels or ROIs to form a time by space matrix, the same voxel or ROI from different subjects was used to create a time by subject matrix. The underlying assumption is that the estimation of neuronal-related activity should not be significantly affected by anatomical and functional differences between subjects once their data has been aligned to a shared space or template. This assumption holds true when analyzing data at the ROI level, where differences are smoothed out during averaging and spatial resolution is reduced. However, when working with voxels, this assumption no longer holds, and differences across subjects are expected. Due to the immense computational cost and memory requirements of simultaneously deconvolving 43 subjects at the voxel level, the analysis in this chapter was performed at the ROI level, where the assumption remained valid. The results demonstrated the capability of the multi-subject PFM (msPFM) algorithm to detect shared and individualized neuronal-related activity across subjects without prior knowledge. Furthermore, the technique successfully linked moment-to-moment brain activity to its underlying causes. Encouragingly, the group synchrony metric showed significant correlations with changes in luminance, audio, speech, the presence of hands, and the presence of faces. Importantly, msPFM was able to adapt to different movies. For example, in the movie "Iteration," where a single character's face was frequently shown, the group synchrony metric correlated with the presence of faces in left middle temporal

gyrus, which is associated with face familiarity and gaze. In contrast, in the movie "Sherlock," where multiple characters appeared and facial identification was more challenging, the group synchrony metric correlated with changes in the presence of faces in the fusiform face areas. These findings indicate that participants may have adapted to certain characteristics of the stimuli, such as luminance and the presence of faces, and that the msPFM estimates were able to capture this adaptation. Although these results are promising, many questions still remain unanswered. For instance, utilizing msPFM could help us understand how the human brain comprehends complex multiplexed signals and identify the specific stimulus features that elicit responses. Furthermore, msPFM has the potential to elucidate the connection between individual differences in these responses and subsequent memory formation or appraisal of the stimulus.

In Chapter 6, the nuclear norm was employed as an additional regularization term in the multivariate PFM formulation. This term effectively tackles the issue of global fluctuations in the BOLD signal, including motion-related signals and physiological artifacts, which can distort the deconvolution of neuronal-related activity. By employing the sparse and low-rank PFM algorithm, the method successfully mitigated this bias, resulting in an accurate estimation of the activity-inducing signal. Notably, the results exhibited a remarkable similarity between the sparse and low-rank PFM algorithm and the single-trial GLM results in detecting neuronal-related activity in a complex dataset characterized by numerous conditions in the experimental task. When the regularization parameter λ was manually selected, the algorithm performed comparably to the multivariate PFM. However, the selection of the regularization parameter for the nuclear norm penalty in the sparse and low-rank PFM posed a significant challenge, which will be addressed in the subsequent section. Given the success of the stability selection procedure in Chapter 3 and Chapter 4, it is worth exploring its adaptation for this algorithm to avoid selecting all three regularization parameters.

Overall, these series of studies demonstrate that PFM techniques can be used to reliably retrieve the neuronal-related activity from fMRI data without any prior information about the experimental paradigm, and that there now exists a formulation of PFM that is suitable for potentially any experimental setting and research question.

Finally, the code used for all the research and algorithms presented in this thesis was written in Python. To promote the use of the PFM techniques developed within this thesis, three separate Python packages have been created and made available as open source software: `pySPFM` for the univariate analysis and as the core library for the other two, `splora` for the multivariate and sparse and low-rank PFM, and `msPFM` for the multi-subject version of PFM. The packages are available on GitHub at https://github.com/paradigm-free-mapping and can be installed using the Python package manager `pip`.

7.2 Future Developments

This thesis showcases research findings that emphasize certain aspects deserving further developments or refinement. First of all, this thesis assumed an identical hemodynamic response model for the entire brain. However, the waveform of the hemodynamic response function (HRF) is known to vary across voxels within cortical regions, across cortical regions, and across subjects (Aguirre, Zarahn, & D'Esposito, 1998b; de Zwart et al., 2005; Handwerker et al., 2004; Miezin et al., 2000; Saad et al., 2001), potentially reflecting different local distributions of vascular anatomy and neurovascular coupling. Therefore, employing prior information about the HRF would allow a more precise estimation of neuronal-related events. Its implementation is straightforward in the case of the univariate formulation of PFM. However, the use of voxel- or region-specific HRFs is not as straightforward in a multivariate scenario. An alternative solution would be to apply the multivariate PFM technique within ROIs using an ROI-specific HRF. An alternative strategy would be to adaptively model the HRF by including the temporal and dispersion derivatives of the assumed canonical HRF in the Toeplitz matrix of the deconvolution model (Gaudes et al., 2012).

Another area of focus to improve PFM is the integration of new regularization terms to mitigate the bias associated with the ℓ_1-norm. In cases where a grouping sparsity constraint is not applied or when dealing with the univariate PFM, an effective method would be to implement regularization with the ℓ_0-norm. However, this presents a challenge as the optimization problem associated with the ℓ_0-norm is known to be a non-convex, NP hard problem. In such cases, an alternative solution could be to utilize $\ell_{0.5}$-norm regularization instead. For the multivariate formulations, a viable option would be to adopt the OSCAR (octagonal selection and clustering algorithm for regression) regularizer. This method involves the use of a combination of ℓ_1 and pair-wise ℓ_∞-norms, which is responsible for its grouping behavior. This approach was proposed to promote group sparsity in situations where the groups are not known beforehand (Bondell & Reich, 2008; Gueddari et al., 2021).

The multivariate formulation of PFM has introduced a significant advancement by enabling the incorporation of spatial information and interactions into the estimation of neuronal-related activity. This opens up possibilities for enhancing the estimation process. For example, the matrix representing the estimated activity-inducing signal $\mathbf{S}$ could be multiplied by a connectivity matrix that represents the interactions between voxels or ROIs in the penalty term. This connectivity matrix could be obtained from other imaging modalities, either structural connectivity from diffusion-weighted MRI or functional connectivity from complementary electrophysiological recordings (e.g., EEG or MEG).

The main limitation of the current multi-subject PFM approach is the assumption that the anatomical and hemodynamic differences between subjects do not significantly affect the estimation of neuronal-related activity when all the data is moved into a shared space. This assumption holds to some extent when working with ROIs, where spatial resolution is reduced and finer anatomical differences across subjects are smoothed out during averaging. However, this

assumption no longer holds striclty in the case of operating at the voxel level. In other words, a perfect anatomical alignment and voxel-to-voxel correspondence in the location of activations cannot be assumed across different subjects. To address the issue of inter-subject spatial variability of functional activations, optimal transport theory could be adopted. This approach, as demonstrated in (Gramfort et al., 2015; Janati et al., 2019, 2020), does not require exact spatial correspondence between neuronal-related events in the group of subjects. Instead, it compares the estimates by considering the geodesic distances between their locations. Furthermore, to tackle the sources of invariance arising from shifts in time, space, and total population size, a more sophisticated formulation that integrates dynamic time warping and unbalanced optimal transport could be considered (Janati et al., 2022).

Furthermore, the selection of the regularization parameter for the nuclear norm in the sparse and low-rank PFM poses a significant challenge. In this thesis, a fixed number of low rank components was chosen as the criterion for selection. However, if functional activations are widespread, this approach may mistakenly classify components resembling BOLD responses as a low-rank, global components. Alternatively, one could consider using a more liberal value for the regularization parameter and then employ a decision tree to distinguish between global and neuronal-related components. For instance, such a decision tree could be developed by drawing inspiration from the ICA AROMA (Pruim, Mennes, van Rooij, et al., 2015) or ICA FIX (Salimi-Khorshidi et al., 2014) methods. Moreover, the application of the low-rank and sparse PFM to resting-state fMRI data remains to be studied.

Future research should also consider the development of deep learning methods for fMRI deconvolution based on physical models of the BOLD signal. Convolutional neural networks (CNNs) have gained significant attention due to their exceptional performance in object classification and segmentation tasks, achieved through training on large image databases. Inspired by these achievements, CNNs have been applied to various inverse problems in imaging, such as denoising, deconvolution, superresolution, and medical image reconstruction (McCann et al., 2017; Wang et al., 2020). These applications have demonstrated promising results, surpassing state-of-the-art techniques, including compressed sensing. Consequently, the implementation of CNNs in PFM techniques has the potential to enhance the accuracy of neuronal-related activity estimates. However, these models are very expensive to train as they require large amounts of data and computational power. Hence, physics-based and self-supervising models could be explored for the development of PFM's deep learning counterparts (Aggarwal et al., 2019; Hammernik et al., 2023; López-Tapia et al., 2021; Lucas et al., 2018).

Finally, the availability of the Python libraries for researchers is crucial. However, in order to promote wider adoption of the PFM techniques, it is essential that these libraries are well documented and include numerous usage examples. Therefore, the development of a comprehensive documentation and tutorial is a priority for the near future.

Appendices

Appendix

A Supplementary Material for Chapter 2

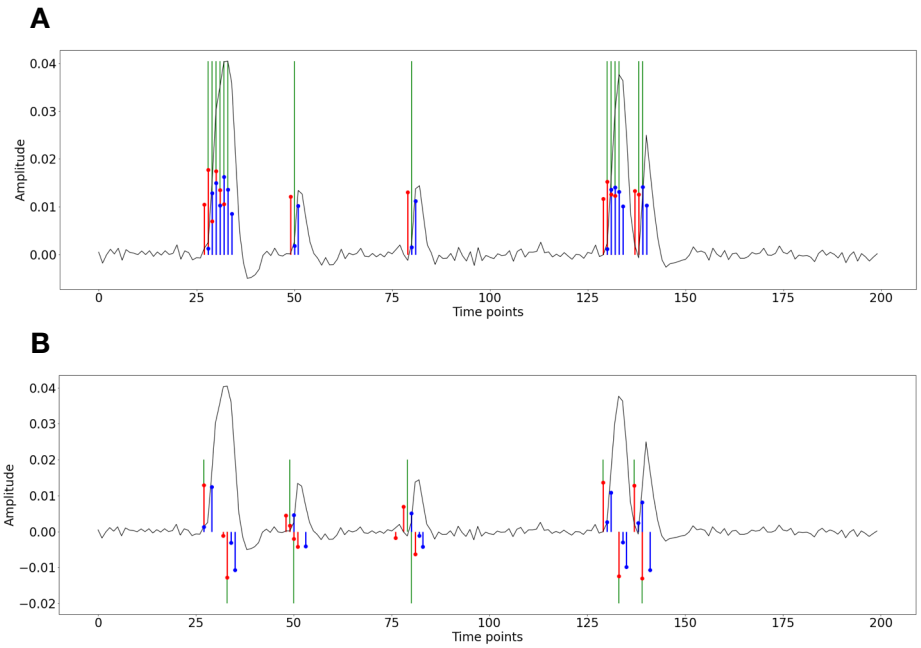

Figure A.1: Activity-inducing (A) and innovation (B) signals estimated with PFM (red) and TA (blue) using their built-in HRF as opposed to using the same. The black line depicts the simulated signal, while the green lines indicate the onsets of the simulated neuronal events. X axis shows time in TRs.

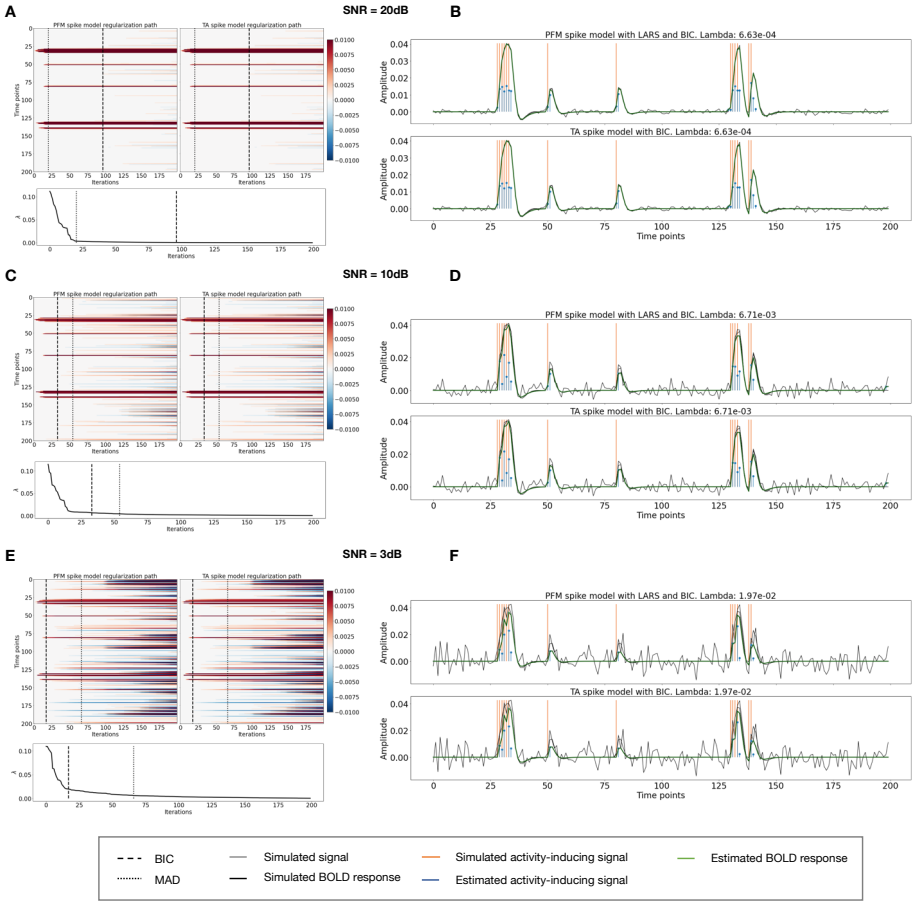

Figure A.2: Spike model simulations. (Left) Heatmap of the regularization paths of the activity-inducing signal estimated with PFM and TA as a function of λ (increasing number of iterations in x-axis), whereas each row in the y-axis shows one time-point. Vertical lines denote iterations corresponding to the Akaike and Bayesian Information Criteria (AIC and BIC) optima. (Right) Estimated activity-inducing (blue) and activity-related (green) signals when set based on BIC. All estimates are identical, regardless of SNR.

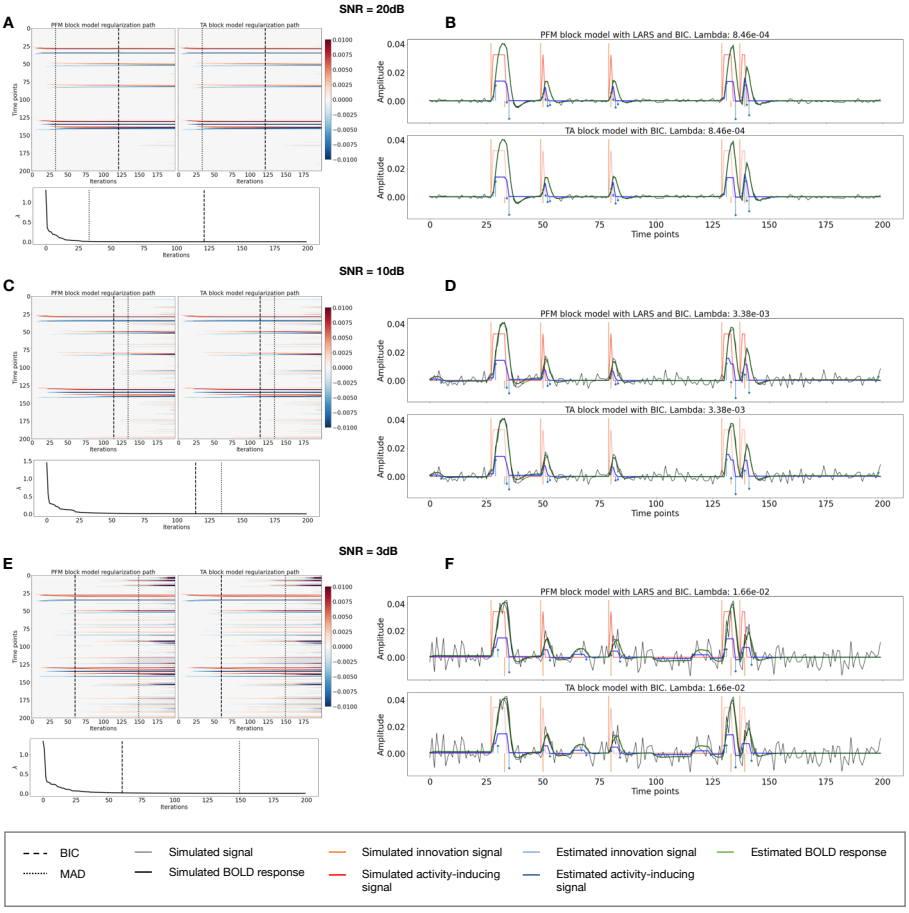

Figure A.3: Block model simulations. (Left) Heatmap of the regularization paths of the innovation signal estimated with PFM and TA as a function of λ (increasing number of iterations in x-axis), whereas each row in the y-axis illustrates one time-point. Vertical lines denote iterations corresponding to the Akaike and Bayesian Information Criteria (AIC and BIC) optima. (Right) Estimated innovation (blue) and activity-related (green) signals when is set based on BIC. All the estimates are identical when compared between the PFM and TA cases, regardless of SNR.

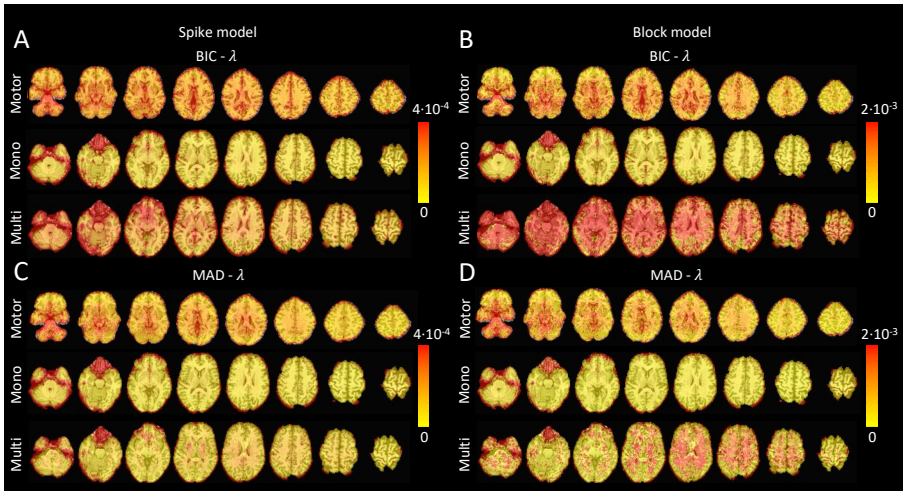

Figure A.4: Values of λ across the different voxels in the brain used to estimate (A) the activity-inducing signal (spike model) and (B) the innovation signal (block model) with the BIC selection, as well as (C) the activity-inducing signal (block model) and (D) the innovation signal (block model) with a MAD-based selection. The λ maps are shown for the three experimental fMRI datasets: the motor task (Motor), the monoband resting-state (Mono), and the multiband resting-state (Multi) datasets.

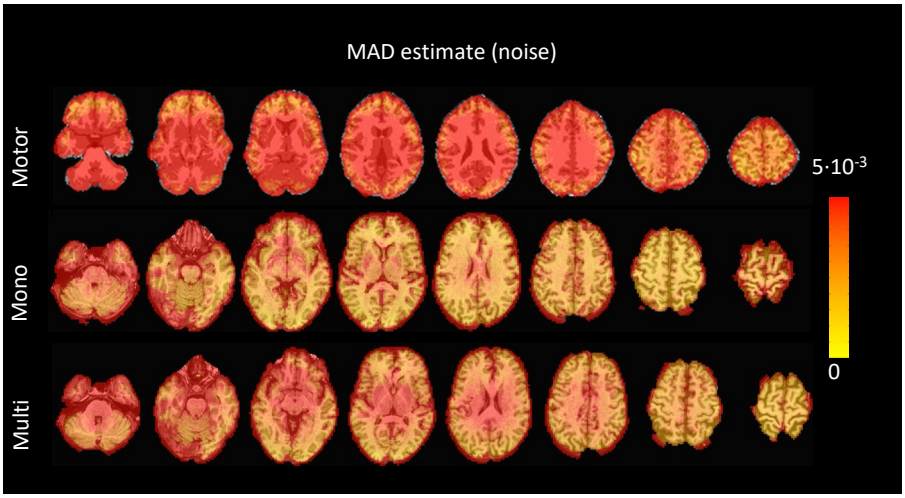

Figure A.5: Values of the MAD estimate of standard deviation of the noise across the different voxels in the brain for the three experimental fMRI datasets: the motor task (Motor), the monoband resting-state (Mono), and the multiband resting-state (Multi) datasets.

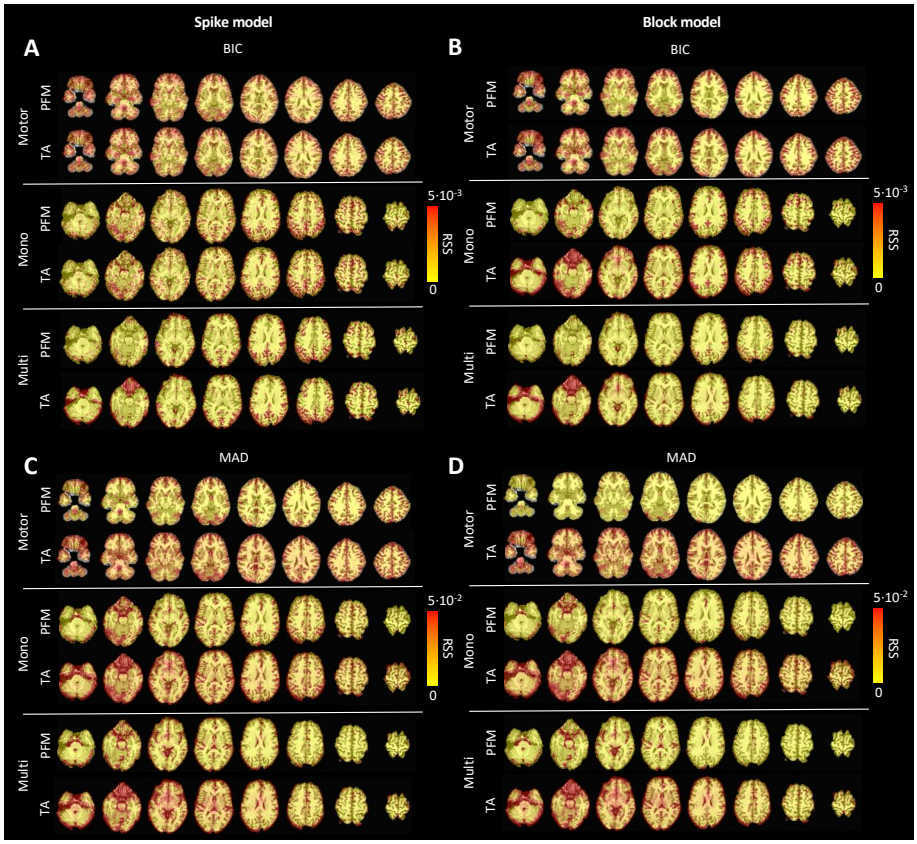

Figure A.6: Root sum of squares (RSS) comparison between Paradigm Free Mapping and Total Activation for the three experimental fMRI datasets: the motor task (Motor), the monoband resting-state (Mono), and the multiband resting-state (Multi) datasets. RSS maps are shown for the spike (left) and block (right) models solved with a selection of λ based on the BIC (top) and MAD (bottom) criteria.

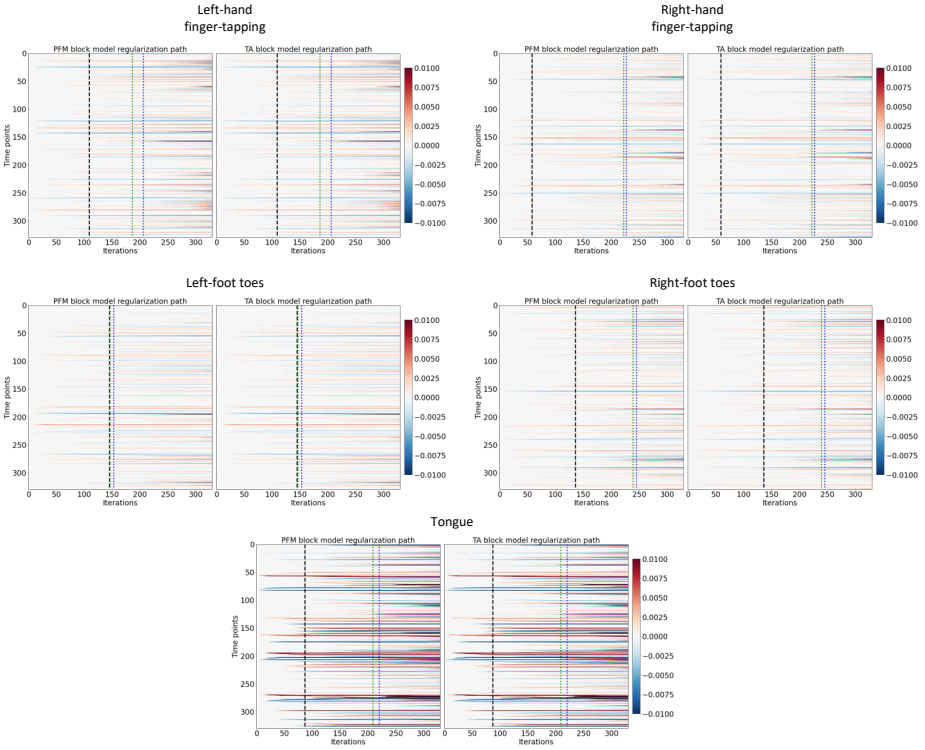

Figure A.7: Regularization paths of the innovation signal estimated with PFM and TA as a function of λ (increasing number of iterations in x-axis, whereas each row in the y-axis shows one time-point) for the representative voxels of the motor task shown in Figure Figure 2.5. Vertical lines denote selections of λ corresponding to the BIC (black), MAD based on LARS residuals (blue) and MAD based on FISTA residuals (green) optima.

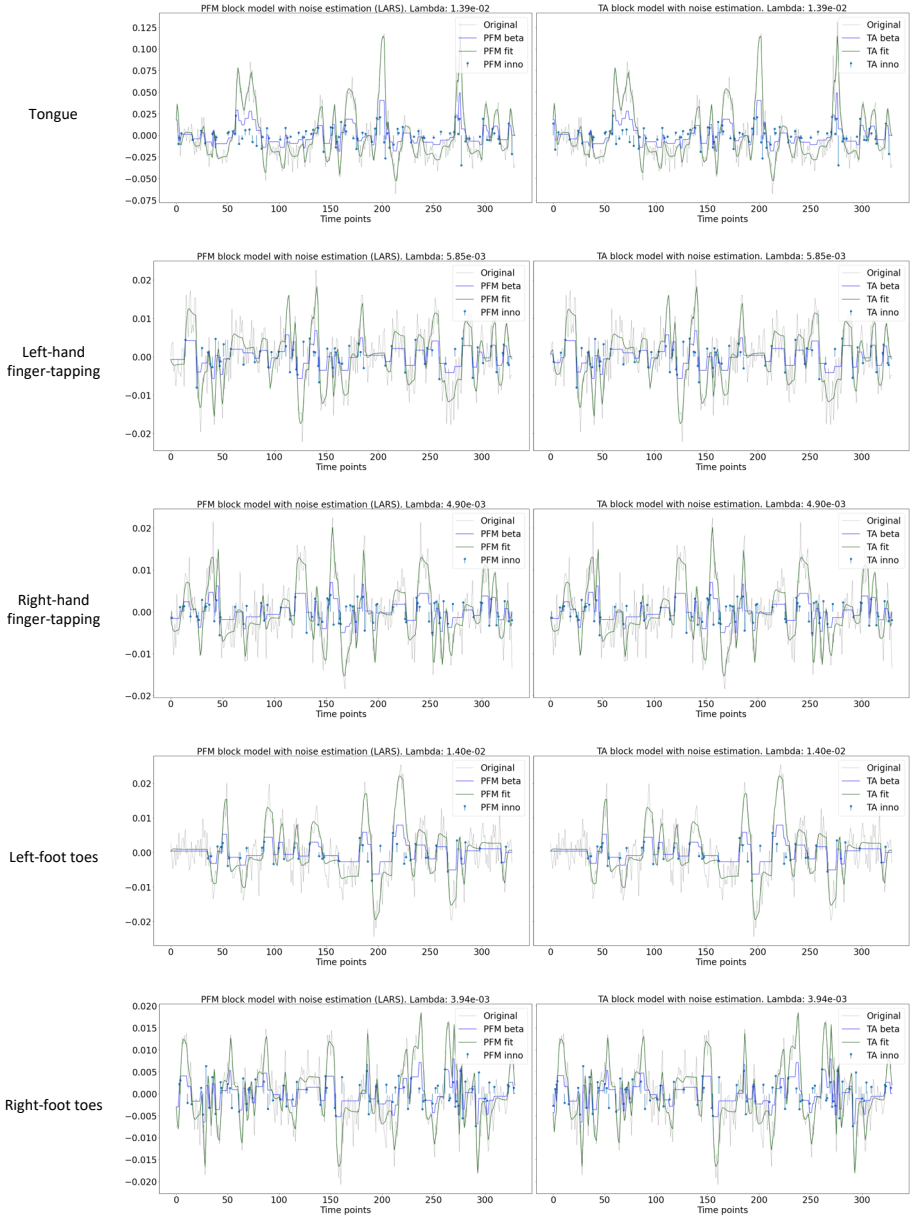

Figure A.8: Estimated innovation signal (blue) and activity-related signal (green) for the representative voxels of the motor task shown in Figure Figure 2.5 with the MAD selection of λ made by TA, i.e., employing the same λ with both PFM and TA.

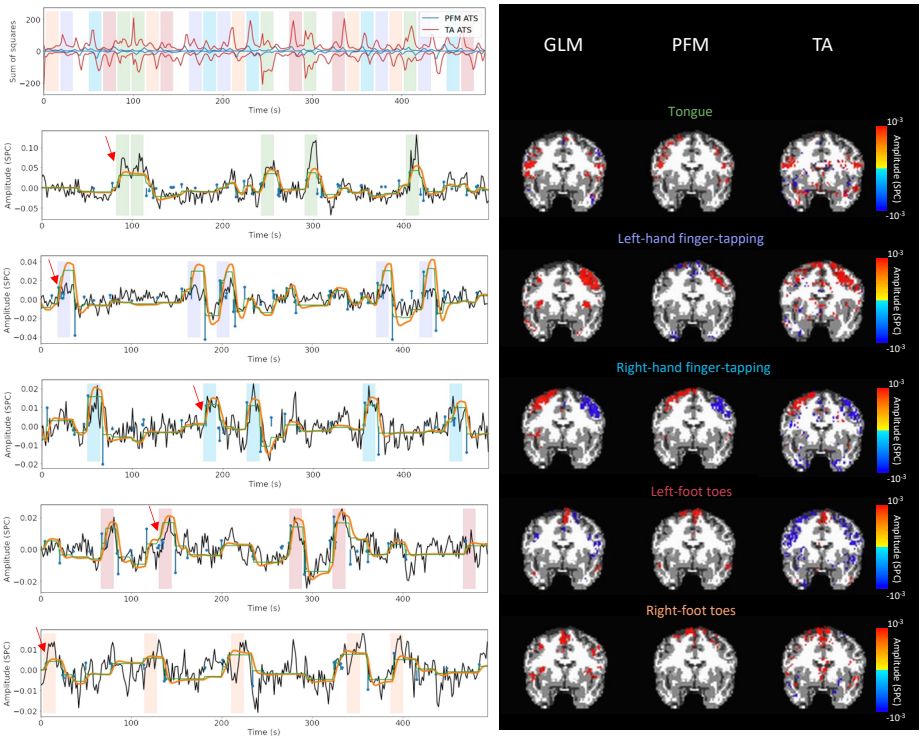

Figure A.9: Activity maps of the motor task using a seletion of λ based on the MAD estimate. Row 1: Activation time-series of the innovation signals estimated by PFM (in blue) or TA (in red) calculated as the sum of squares of all voxels at every timepoint. Positive-valued and negative-valued contributions were separated into two distinct time-courses. Color-bands indicate the onset and duration of each condition in the task (green: tongue, purple: left-hand finger-tapping, blue: right-hand finger-tapping, red: left-foot toes, orange: right-foot toes). Rows 2-6: time-series of a representative voxel for each task with the PFM-estimated innovation (blue), PFM-estimated activity-inducing (green), and activity-related (i.e., fitted, orange) signals, with their corresponding GLM, PFM, and TA maps on the right. The maps shown on the right are sampled at the time-point labeled with the red arrows and display the innovation signals at that moment across the whole brain.

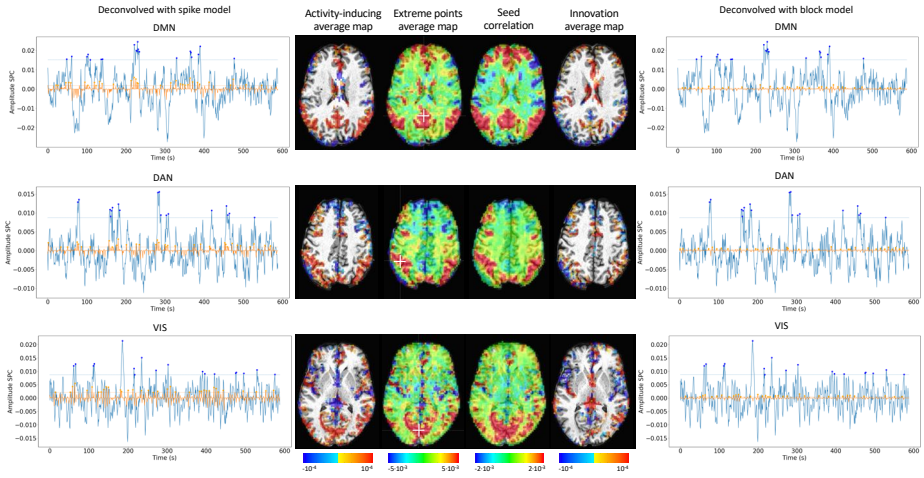

Figure A.10: Activity-inducing CAPs (left) and innovation CAPs (right) obtained with the PFM-estimated activity-inducing and innovation signals respectively, using a MAD-based selection of λ. Time-points selected with a 95th percentile threshold are shown over the average time-series (blue) in the seed region (white-cross) and the deconvolved signal (orange). CAPs and seed correlation maps are illustrated in the center.

Bibliography

Abbas, A., Belloy, M., Kashyap, A., Billings, J., Nezafati, M., Schumacher, E. H., & Keilholz, S. (2019). Quasi-periodic patterns contribute to functional connectivity in the brain. NeuroImage, vol. 191, 193–204.

Adelstein, J. S., Shehzad, Z., Mennes, M., DeYoung, C. G., Zuo, X.-N., Kelly, C., Margulies, D. S., Bloomfield, A., Gray, J. R., Castellanos, F. X., & Milham, M. P. (2011). Personality is reflected in the brain's intrinsic functional architecture (M. Valdes-Sosa, Ed.). PLoS ONE, vol. 6no. 11, e27633.

Aggarwal, H. K., Mani, M. P., & Jacob, M. (2019). Modl: Model-based deep learning architecture for inverse problems. IEEE Transactions on Medical Imaging, vol. 38no. 2, 394–405.

Aguirre, G. K., Zarahn, E., & D'Esposito, M. (1998a). An area within human ventral cortex sensitive to "building" stimuli. Neuron, vol. 21no. 2, 373–383.

Aguirre, G., Zarahn, E., & D'Esposito, M. (1998b). The variability of human, BOLD hemodynamic responses. NeuroImage, vol. 8no. 4, 360–369.

Akçakaya, M., Yaman, B., Chung, H., & Ye, J. C. (2021). Unsupervised deep learning methods for biological image reconstruction. arXiv.

Allan, T. W., Francis, S. T., Caballero-Gaudes, C., Morris, P. G., Liddle, E. B., Liddle, P. F., Brookes, M. J., & Gowland, P. A. (2015). Functional connectivity in MRI is driven by spontaneous BOLD events (E. A. Stamatakis, Ed.). PLOS ONE, vol. 10no. 4, e0124577.

Allen, E. A., Damaraju, E., Plis, S. M., Erhardt, E. B., Eichele, T., & Calhoun, V. D. (2012). Tracking whole-brain connectivity dynamics in the resting state. Cerebral Cortex, vol. 24no. 3, 663–676.

Amemiya, S., Yamashita, H., Takao, H., & Abe, O. (2018). Integrated multi-echo denoising strategy improves identification of inherent language laterality. Magnetic Resonance in Medicine, vol. 81no. 5, 3262–3271.

Amico, E., & Goñi, J. (2018). The quest for identifiability in human functional connectomes. Scientific Reports, vol. 8no. 1.

Anderson, J. S., Druzgal, T. J., Lopez-Larson, M., Jeong, E.-K., Desai, K., & Yurgelun-Todd, D. (2011). Network anticorrelations, global regression, and phase-shifted soft tissue correction. Human Brain Mapping, vol. 32no. 6, 919–934.

Andersson, J. L., Skare, S., & Ashburner, J. (2003). How to correct susceptibility distortions in spin-echo echo-planar images: Application to diffusion tensor imaging. NeuroImage, vol. 20no. 2, 870–888.

Arieli, A., Shoham, D., Hildesheim, R., & Grinvald, A. (1995). Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. Journal of Neurophysiology, vol. 73no. 5, 2072–2093.

Aslan, S., Cemgil, A. T., & Akın, A. (2016). Joint state and parameter estimation of the hemodynamic model by particle smoother expectation maximization method. Journal of Neural Engineering, vol. 13no. 4, 046010.

Attarpour, A., Ward, J., & Chen, J. J. (2021). Vascular origins of low-frequency oscillations in the cerebrospinal fluid signal in resting-state fmri: Interpretation using photoplethysmography. Human Brain Mapping, vol. 42no. 8, 2606–2622.

Atwi, S., Metcalfe, A. W. S., Robertson, A. D., Rezmovitz, J., Anderson, N. D., & MacIntosh, B. J. (2018). Attention-related brain activation is altered in older adults with white matter hyperintensities using multi-echo fMRI. Frontiers in Neuroscience, vol. 12.

Avidan, G., & Behrmann, M. (2010). Cortical networks mediating face familiarity and identity in the human brain. Journal of Vision, vol. 5no. 8, 633–633.

Bailey, H. R., Kurby, C. A., Sargent, J. Q., & Zacks, J. M. (2017). Attentional focus affects how events are segmented and updated in narrative reading. Memory & Cognition, vol. 45no. 6, 940–955.

Baillet, S., Mosher, J., & Leahy, R. (2001). Electromagnetic brain mapping. IEEE Signal Processing Magazine, vol. 18no. 6, 14–30.

Baldassano, C., Chen, J., Zadbood, A., Pillow, J. W., Hasson, U., & Norman, K. A. (2017). Discovering event structure in continuous narrative perception and memory. Neuron, vol. 95no. 3, 709–721.e5.

Bandettini, P. A., Wong, E. C., Hinks, R. S., Tikofsky, R. S., & Hyde, J. S. (1992). Time course EPI of human brain function during task activation. Magnetic Resonance in Medicine, vol. 25no. 2, 390–397.

Baseler, H. A., Harris, R. J., Young, A. W., & Andrews, T. J. (2012). Neural responses to expression and gaze in the posterior superior temporal sulcus interact with facial identity. Cerebral Cortex, vol. 24no. 3, 737–744.

Beall, E. B., & Lowe, M. J. (2007). Isolating physiologic noise sources with independently determined spatial measures. NeuroImage, vol. 37no. 4, 1286–1300.

Beall, E. B., & Lowe, M. J. (2014). Simpace: Generating simulated motion corrupted bold data with synthetic-navigated acquisition for the development and evaluation of slomoco: A new, highly effective slicewise motion correction. NeuroImage, vol. 101, 21–34.

Beck, A., & Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM Journal on Imaging Sciences, vol. 2no. 1, 183–202.

Behzadi, Y., & Liu, T. T. (2006). Caffeine reduces the initial dip in the visual BOLD response at 3 t. NeuroImage, vol. 32no. 1, 9–15.

Behzadi, Y., Restom, K., Liau, J., & Liu, T. T. (2007). A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. NeuroImage, vol. 37no. 1, 90–101.

Belliveau, J. W., Kennedy, D. N., McKinstry, R. C., Buchbinder, B. R., Weisskoff, R. M., Cohen, M. S., Vevea, J. M., Brady, T. J., & Rosen, B. R. (1991). Functional mapping of the human visual cortex by magnetic resonance imaging. Science, vol. 254no. 5032, 716–719.

Bernhardt, B. C., Smallwood, J., Keilholz, S., & Margulies, D. S. (2022). Gradients in brain organization. NeuroImage, vol. 251, 118987.

Bianciardi, M., Fukunaga, M., van Gelderen, P., Horovitz, S. G., de Zwart, J. A., Shmueli, K., & Duyn, J. H. (2009). Sources of functional magnetic resonance imaging signal fluctuations in the human brain at rest: A 7 t study. Magnetic Resonance Imaging, vol. 27no. 8, 1019–1029.

Birn, R. M., Diamond, J. B., Smith, M. A., & Bandettini, P. A. (2006). Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. NeuroImage, vol. 31no. 4, 1536–1548.

Birn, R. M., Smith, M. A., Jones, T. B., & Bandettini, P. A. (2008). The respiration response function: The temporal dynamics of fMRI signal fluctuations related to changes in respiration. NeuroImage, vol. 40no. 2, 644–654.

Biswal, B., Deyoe, E. A., & Hyde, J. S. (1996). Reduction of physiological fluctuations in fMRI using digital filters. Magnetic Resonance in Medicine, vol. 35no. 1, 107–113.

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, vol. 34no. 4, 537–541.

Bolt, T., Nomi, J. S., Bzdok, D., Salas, J. A., Chang, C., Yeo, B. T. T., Uddin, L. Q., & Keilholz, S. D. (2022). A parsimonious description of global functional brain organization in three spatiotemporal patterns. Nature Neuroscience, vol. 25no. 8, 1093–1103.

Bolton, T. A., Farouj, Y., Inan, M., & Ville, D. V. D. (2019). Structurally-informed deconvolution of functional magnetic resonance imaging data. 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019).

Bolton, T. A., Freitas, L. G., Jochaut, D., Giraud, A.-L., & Ville, D. V. D. (2020). Neural responses in autism during movie watching: Inter-individual re-

sponse variability co-varies with symptomatology. NeuroImage, vol. 216, 116571.

Bolton, T. A., Jochaut, D., Giraud, A.-L., & Ville, D. V. D. (2018). Brain dynamics in ASD during movie-watching show idiosyncratic functional integration and segregation. Human Brain Mapping, vol. 39no. 6, 2391–2404.

Bolton, T. A., Morgenroth, E., Preti, M. G., & Van De Ville, D. (2020). Tapping into multi-faceted human behavior and psychopathology using fMRI brain dynamics. Trends in Neurosciences, vol. 43no. 9, 667–680.

Bommarito, G., Tarun, A., Farouj, Y., Preti, M. G., Petracca, M., Droby, A., El Mendili, M. M., Inglese, M., & Van De Ville, D. (2021). Altered anterior default mode network dynamics in progressive multiple sclerosis. Multiple Sclerosis Journal, vol. 28no. 2, 206–216.

Bondell, H. D., & Reich, B. J. (2008). Simultaneous regression shrinkage, variable selection, and supervised clustering of predictors with OSCAR. Biometrics, vol. 64no. 1, 115–123.

Boynton, G. M., Engel, S. A., Glover, G. H., & Heeger, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human v1. The Journal of Neuroscience, vol. 16no. 13, 4207–4221.

Bright, M. G., & Murphy, K. (2013). Removing motion and physiological artifacts from intrinsic BOLD fluctuations using short echo data. NeuroImage, vol. 64, 526–537.

Bruckstein, A. M., Donoho, D. L., & Elad, M. (2009). From sparse solutions of systems of equations to sparse modeling of signals and images. SIAM Review, vol. 51no. 1, 34–81.

Buckner, R., Bandettini, P., O'Craven, K., Savoy, R., Petersen, S., Raichle, M., & Rosen, B. (1996). Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic resonance imaging. Proceedings of the National Academy of Sciences, vol. 93no. 25, 14878–14883.

Bulte, D., Chiarelli, P., Wise, R., & Jezzard, P. (2007). Measurement of cerebral blood volume in humans using hyperoxic mri contrast. Journal of Magnetic Resonance Imaging, vol. 26no. 4, 894–899.

Bush, K., & Cisler, J. (2013). Decoding neural events from fMRI BOLD signal: A comparison of existing approaches and development of a new algorithm. Magnetic Resonance Imaging, vol. 31no. 6, 976–989.

Bush, K., Cisler, J., Bian, J., Hazaroglu, G., Hazaroglu, O., & Kilts, C. (2015). Improving the precision of fMRI BOLD signal deconvolution with implications for connectivity analysis. Magnetic Resonance Imaging, vol. 33no. 10, 1314–1323.

Bush, K., Zhou, S., Cisler, J., Bian, J., Hazaroglu, O., Gillispie, K., Yoshigoe, K., & Kilts, C. (2015). A deconvolution-based approach to identifying large-scale effective connectivity. Magnetic Resonance Imaging, vol. 33no. 10, 1290–1298.

Buxton, R. B. (2009, August). Introduction to functional magnetic resonance imaging. Cambridge University Press.

Buxton, R. B., Wong, E. C., & Frank, L. R. (1998). Dynamics of blood flow and oxygenation changes during brain activation: The balloon model. Magnetic resonance in medicine, vol. 39no. 6, 855–864.

Caballero-Gaudes, C., Moia, S., Panwar, P., Bandettini, P. A., & Gonzalez-Castillo, J. (2019). A deconvolution algorithm for multi-echo functional MRI: Multi-echo sparse paradigm free mapping. NeuroImage, vol. 202, 116081.

Caballero-Gaudes, C., & Reynolds, R. C. (2017). Methods for cleaning the BOLD fMRI signal. NeuroImage, vol. 154, 128–149.

Cabral, J., Vidaurre, D., Marques, P., Magalhães, R., Moreira, P. S., Soares, J. M., Deco, G., Sousa, N., & Kringelbach, M. L. (2017). Cognitive performance in healthy older adults relates to spontaneous switching between states of functional connectivity during rest. Scientific Reports, vol. 7no. 1.

Casanova, R., Ryali, S., Serences, J., Yang, L., Kraft, R., Laurienti, P. J., & Maldjian, J. A. (2008). The impact of temporal regularization on estimates of the BOLD hemodynamic response function: A comparative analysis. NeuroImage, vol. 40no. 4, 1606–1618.

Chambolle, A. (2004). An algorithm for total variation minimization and applications. Journal of Mathematical imaging and vision, vol. 20no. 1, 89–97.

Chang, C., Cunningham, J. P., & Glover, G. H. (2009). Influence of heart rate on the BOLD signal: The cardiac response function. NeuroImage, vol. 44no. 3, 857–869.

Chang, C., & Glover, G. H. (2009). Relationship between respiration, end-tidal CO2, and BOLD signals in resting-state fMRI. NeuroImage, vol. 47no. 4, 1381–1393.

Chang, C., & Glover, G. H. (2010). Time–frequency dynamics of resting-state brain connectivity measured with fMRI. NeuroImage, vol. 50no. 1, 81–98.

Chen, G., Taylor, P. A., Qu, X., Molfese, P. J., Bandettini, P. A., Cox, R. W., & Finn, E. S. (2020). Untangling the relatedness among correlations, part iii: Inter-subject correlation analysis through bayesian multilevel modeling for naturalistic scanning. NeuroImage, vol. 216, 116474.

Chen, J., Leong, Y. C., Honey, C. J., Yong, C. H., Norman, K. A., & Hasson, U. (2017). Shared memories reveal shared structure in neural activity across individuals. Nature Neuroscience, vol. 20no. 1, 115–125.

Chen, J. J., & Pike, G. B. (2009). Origins of the BOLD post-stimulus undershoot. NeuroImage, vol. 46no. 3, 559–568.

Chen, J. E., Chang, C., Greicius, M. D., & Glover, G. H. (2015). Introducing co-activation pattern metrics to quantify spontaneous brain network dynamics. NeuroImage, vol. 111, 476–488.

Chen, J. E., Glover, G. H., Fultz, N. E., Rosen, B. R., Polimeni, J. R., & Lewis, L. D. (2021). Investigating mechanisms of fast BOLD responses: The effects of stimulus intensity and of spatial heterogeneity of hemodynamics. NeuroImage, vol. 245, 118658.

Chen, P.-H., Chen, J., Yeshurun, Y., Hasson, U., Haxby, J., & Ramadge, P. J. (2015). A reduced-dimension fmri shared response model. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, & R. Garnett (Eds.), Advances in neural information processing systems (Vol. 28). Curran Associates, Inc.

Chen, S. S., Donoho, D. L., & Saunders, M. A. (2001). Atomic decomposition by basis pursuit. SIAM review, vol. 43no. 1, 129–159.

Cherkaoui, H., Moreau, T., Halimi, A., & Ciuciu, P. (2019). Sparsity-based blind deconvolution of neural activation signal in fMRI. ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 1323–1327.

Cherkaoui, H., Moreau, T., Halimi, A., Leroy, C., & Ciuciu, P. (2021). Multivariate semi-blind deconvolution of fMRI time series. NeuroImage, vol. 241, 118418.

Cherkaoui, H., Sulam, J., & Moreau, T. (2020). Learning to solve TV regularised problems with unrolled algorithms (H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin, Eds.). Advances in Neural Information Processing Systems, vol. 33, 11513–11524.

Christen, T., Ni, W., Qiu, D., Schmiedeskamp, H., Bammer, R., Moseley, M., & Zaharchuk, G. (2012). High-resolution cerebral blood volume imaging in humans using the blood pool contrast agent ferumoxytol: High resolution cbv maps in humans using uspios. Magnetic Resonance in Medicine, vol. 70no. 3, 705–710.

Chuang, K.-H., & Chen, J.-H. (2001). IMPACT: Image-based physiological artifacts estimation and correction technique for functional MRI. Magnetic Resonance in Medicine, vol. 46no. 2, 344–353.

Cifre, I., Zarepour, M., Horovitz, S. G., Cannas, S. A., & Chialvo, D. R. (2020). Further results on why a point process is effective for estimating correlation between brain regions. Papers in Physics, vol. 12, 120003.

Cifre, I., Flores, M. T. M., Ochab, J. K., & Chialvo, D. R. (2020). Revisiting non-linear functional brain co-activations: Directed, dynamic and delayed.

Ciuciu, P., Poline, J.-B., Marrelec, G., Idier, J., Pallier, C., & Benali, H. (2003). Unsupervised robust nonparametric estimation of the hemodynamic response function for any fmri experiment. IEEE transactions on medical imaging, vol. 22, 1235–1251.

Cohen, A. D., Jagra, A. S., Visser, N. J., Yang, B., Fernandez, B., Banerjee, S., & Wang, Y. (2021). Improving the breath-holding CVR measurement using the multiband multi-echo EPI sequence. Frontiers in Physiology, vol. 12.

Cohen, A. D., & Wang, Y. (2019). Improving the assessment of breath-holding induced cerebral vascular reactivity using a multiband multi-echo ASL/BOLD sequence. Scientific Reports, vol. 9no. 1.

Cohen, M. S. (1997). Parametric analysis of fMRI data using linear systems methods. NeuroImage, vol. 6no. 2, 93–103.

Cole, M. W., Yarkoni, T., Repovs, G., Anticevic, A., & Braver, T. S. (2012). Global connectivity of prefrontal cortex predicts cognitive control and intelligence. Journal of Neuroscience, vol. 32no. 26, 8988–8999.

Costantini, I., Deriche, R., & Deslauriers-Gauthier, S. (2022). An anisotropic 4d filtering approach to recover brain activation from paradigm-free functional MRI data. Frontiers in Neuroimaging, vol. 1.

Cox, R. W. (1996). AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. Computers and Biomedical Research, vol. 29no. 3, 162–173.

Cribben, I., Wager, T. D., & Lindquist, M. A. (2013). Detecting functional connectivity change points for single-subject fmri data. Frontiers in Computational Neuroscience, vol. 7.

Dale, A. M., & Halgren, E. (2001). Spatiotemporal mapping of brain activity by integration of multiple imaging modalities. Current Opinion in Neurobiology, vol. 11no. 2, 202–208.

de la Vega, A., Rocca, R., Blair, R. W., Markiewicz, C. J., Mentch, J., Kent, J. D., Herholz, P., Ghosh, S. S., Poldrack, R. A., & Yarkoni, T. (2022). Neuroscout, a unified platform for generalizable and reproducible fMRI research. eLife, vol. 11.

Dean, B. L., Lee, C., Kirsch, J. E., Runge, V. M., Dempsey, R. M., Luther, & Pettigrew. (1992). Cerebral hemodynamics and cerebral blood volume: Mr assessment using gadolinium contrast agents and t1-weighted turbo-flash imaging. AJNR. American journal of neuroradiology, vol. 13 1, 39–48.

Deco, G., & Kringelbach, M. L. (2017). Hierarchy of information processing in the brain: A novel 'intrinsic ignition' framework. Neuron, vol. 94, 961–968.

Deco, G., Tagliazucchi, E., Laufs, H., Sanjuán, A., & Kringelbach, M. L. (2017). Novel intrinsic ignition method measuring local-global integration characterizes wakefulness and deep sleep. eNeuro, vol. 4.

de Zwart, J. A., Silva, A. C., van Gelderen, P., Kellman, P., Fukunaga, M., Chu, R., Koretsky, A. P., Frank, J. A., & Duyn, J. H. (2005). Temporal dynamics of the BOLD fMRI impulse response. NeuroImage, vol. 24no. 3, 667–677.

Di, X., & Biswal, B. B. (2013). Modulatory interactions of resting-state brain functional connectivity (Q. Gong, Ed.). PLoS ONE, vol. 8no. 8, e71163.

Di, X., & Biswal, B. B. (2015). Characterizations of resting-state modulatory interactions in the human brain. Journal of Neurophysiology, vol. 114no. 5, 2785–2796.

Di, X., & Biswal, B. B. (2018). Toward task connectomics: Examining whole-brain task modulated connectivity in different task domains. Cerebral Cortex, vol. 29no. 4, 1572–1583.

Di, X., & Biswal, B. B. (2020). Intersubject consistent dynamic connectivity during natural vision revealed by functional MRI. NeuroImage, vol. 216, 116698.

Dipasquale, O., Sethi, A., Laganà, M. M., Baglio, F., Baselli, G., Kundu, P., Harrison, N. A., & Cercignani, M. (2017). Comparing resting state fMRI de-noising approaches using multi- and single-echo acquisitions (A. Gozzi, Ed.). PLOS ONE, vol. 12no. 3, e0173289.

Dosenbach, N. U. F., Nardos, B., Cohen, A. L., Fair, D. A., Power, J. D., Church, J. A., Nelson, S. M., Wig, G. S., Vogel, A. C., Lessov-Schlaggar, C. N., Barnes, K. A., Dubis, J. W., Feczko, E., Coalson, R. S., Pruett, J. R., Barch, D. M., Petersen, S. E., & Schlaggar, B. L. (2010). Prediction of individual brain maturity using fMRI. Science, vol. 329no. 5997, 1358–1361.

Duann, J.-R., Jung, T.-P., Kuo, W.-J., Yeh, T.-C., Makeig, S., Hsieh, J.-C., & Sejnowski, T. J. (2002). Single-trial variability in event-related BOLD signals. NeuroImage, vol. 15no. 4, 823–835.

Dubois, J., Galdi, P., Han, Y., Paul, L. K., & Adolphs, R. (2018). Resting-state functional brain connectivity best predicts the personality dimension of openness to experience. Personality Neuroscience, vol. 1.

Duong, T. Q., Kim, D.-S., U?urbil, K., & Kim, S.-G. (2000). Spatiotemporal dynamics of the BOLD fMRI signals: Toward mapping submillimeter cortical columns using the early negative response. Magnetic Resonance in Medicine, vol. 44no. 2, 231–242.

DuPre, E., Salo, T., Ahmed, Z., Bandettini, P., Bottenhorn, K., Caballero-Gaudes, C., Dowdle, L., Gonzalez-Castillo, J., Heunis, S., Kundu, P., Laird, A., Markello, R., Markiewicz, C., Moia, S., Staden, I., Teves, J., Uruñuela, E., Vaziri-Pashkam, M., Whitaker, K., & Handwerker, D. (2021). TE-dependent analysis of multi-echo fMRI with tedana. Journal of Open Source Software, vol. 6no. 66, 3669.

Efron, B., Hastie, T., Johnstone, I., & Tibshirani, R. (2004). Least angle regression. The Annals of Statistics, vol. 32no. 2, 407–499.

Eger, E., Schweinberger, S., Dolan, R., & Henson, R. (2005). Familiarity enhances invariance of face representations in human ventral visual cortex: Fmri evidence. NeuroImage, vol. 26no. 4, 1128–1139.

Elad, M., Milanfar, P., & Rubinstein, R. (2007). Analysis versus synthesis in signal priors. Inverse Problems, vol. 23no. 3, 947–968.

Emerson, R. W., Adams, C., Nishino, T., Hazlett, H. C., Wolff, J. J., Zwaigenbaum, L., Constantino, J. N., Shen, M. D., Swanson, M. R., Elison, J. T., Kandala, S., Estes, A. M., Botteron, K. N., Collins, L., Dager, S. R., Evans, A. C., Gerig, G., Gu, H., McKinstry, R. C., ... and, J. P. (2017). Functional neuroimaging of high-risk 6-month-old infants predicts a diagnosis of autism at 24 months of age. Science Translational Medicine, vol. 9no. 393.

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. Nature, vol. 392no. 6676, 598–601.

Ernst, T., & Hennig, J. (1994). Observation of a fast response in functional MR. Magnetic Resonance in Medicine, vol. 32no. 1, 146–149.

Esfahlani, F. Z., Jo, Y., Faskowitz, J., Byrge, L., Kennedy, D. P., Sporns, O., & Betzel, R. F. (2020). High-amplitude cofluctuations in cortical activity drive functional connectivity. Proceedings of the National Academy of Sciences, vol. 117no. 45, 28393–28401.

Esteban, O., Markiewicz, C. J., Blair, R. W., Moodie, C. A., Isik, A. I., Erramuzpe, A., Kent, J. D., Goncalves, M., DuPre, E., Snyder, M., Oya, H., Ghosh,

S. S., Wright, J., Durnez, J., Poldrack, R. A., & Gorgolewski, K. J. (2018). fMRIPrep: A robust preprocessing pipeline for functional MRI. Nature Methods, vol. 16no. 1, 111–116.

Evans, J. W., Kundu, P., Horovitz, S. G., & Bandettini, P. A. (2015). Separating slow BOLD from non-BOLD baseline drifts using multi-echo fMRI. NeuroImage, vol. 105, 189–197.

Fair, D. A., Miranda-Dominguez, O., Snyder, A. Z., Perrone, A., Earl, E. A., Van, A. N., Koller, J. M., Feczko, E., Tisdall, M. D., van der Kouwe, A., Klein, R. L., Mirro, A. E., Hampton, J. M., Adeyemo, B., Laumann, T. O., Gratton, C., Greene, D. J., Schlaggar, B. L., Hagler, D. J., … Dosenbach, N. U. (2020). Correction of respiratory artifacts in MRI head motion estimates. NeuroImage, vol. 208, 116400.

Farouj, Y., Karahanoglu, F. I., & Van De Ville, D. (2019). Bold signal deconvolution under uncertain hæmodynamics: A semi-blind approach. 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), 1792–1796.

Farouj, Y., Karahanoglu, F. I., & Ville, D. V. D. (2017). Regularized spatiotemporal deconvolution of fMRI data using gray-matter constrained total variation. 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017).

Farouj, Y., Karahanoglu, F. I., & Van De Ville, D. (2020). Deconvolution of sustained neural activity from large-scale calcium imaging data. IEEE Transactions on Medical Imaging, vol. 39no. 4, 1094–1103.

Faskowitz, J., Esfahlani, F. Z., Jo, Y., Sporns, O., & Betzel, R. F. (2020). Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture. Nature Neuroscience, vol. 23no. 12, 1644–1654.

Feinberg, D. A., Moeller, S., Smith, S. M., Auerbach, E., Ramanna, S., Glasser, M. F., Miller, K. L., Ugurbil, K., & Yacoub, E. (2010). Multiplexed echo planar imaging for sub-second whole brain FMRI and fast diffusion imaging (P. A. Valdes-Sosa, Ed.). PLoS ONE, vol. 5no. 12, e15710.

Fernandez, B., Leuchs, L., Sämann, P. G., Czisch, M., & Spoormaker, V. I. (2017). Multi-echo EPI of human fear conditioning reveals improved BOLD detection in ventromedial prefrontal cortex. NeuroImage, vol. 156, 65–77.

Finn, E. S. (2021). Is it time to put rest to rest? Trends in Cognitive Sciences, vol. 25no. 12, 1021–1032.

Finn, E. S., & Bandettini, P. A. (2021). Movie-watching outperforms rest for functional connectivity-based prediction of behavior. NeuroImage, vol. 235, 117963.

Finn, E. S., Glerean, E., Khojandi, A. Y., Nielson, D., Molfese, P. J., Handwerker, D. A., & Bandettini, P. A. (2020). Idiosynchrony: From shared responses to individual differences during naturalistic neuroimaging. NeuroImage, vol. 215, 116828.

Finn, E. S., Scheinost, D., Finn, D. M., Shen, X., Papademetris, X., & Constable, R. T. (2017). Can brain state be manipulated to emphasize individual differences in functional connectivity? NeuroImage, vol. 160, 140–151.

Finn, E. S., Shen, X., Scheinost, D., Rosenberg, M. D., Huang, J., Chun, M. M., Papademetris, X., & Constable, R. T. (2015). Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. Nature Neuroscience, vol. 18no. 11, 1664–1671.

Fong, A. H. C., Yoo, K., Rosenberg, M. D., Zhang, S., Li, C.-S. R., Scheinost, D., Constable, R. T., & Chun, M. M. (2019). Dynamic functional connectivity during task performance and rest predicts individual differences in attention across studies. NeuroImage, vol. 188, 14–25.

Fox, P. T., Mintun, M. A., Raichle, M. E., Miezin, F. M., Allman, J. M., & Essen, D. C. V. (1986). Mapping human visual cortex with positron emission tomography. Nature, vol. 323no. 6091, 806–809.

Freitas, L. G., Bolton, T. A., Krikler, B. E., Jochaut, D., Giraud, A.-L., Hüppi, P. S., & Van De Ville, D. (2020). Time-resolved effective connectivity in task fMRI: Psychophysiological interactions of co-activation patterns. NeuroImage, vol. 212, 116635.

Friedrich, J., Zhou, P., & Paninski, L. (2017). Fast online deconvolution of calcium imaging data (J. Vogelstein, Ed.). PLOS Computational Biology, vol. 13no. 3, e1005423.

Friston, K. (2007). Statistical parametric mapping. In Statistical parametric mapping (pp. 10–31). Elsevier.

Friston, K. J., Frith, C. D., Liddle, P. F., & Frackowiak, R. S. J. (1993). Functional connectivity: The principal-component analysis of large (PET) data sets. Journal of Cerebral Blood Flow & Metabolism, vol. 13no. 1, 5–14.

Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J.-P., Frith, C. D., & Frackowiak, R. S. J. (1994). Statistical parametric maps in functional imaging: A general linear approach. Human Brain Mapping, vol. 2no. 4, 189–210.

Friston, K. J., Jezzard, P., & Turner, R. (1994). Analysis of functional MRI time-series. Human Brain Mapping, vol. 1no. 2, 153–171.

Friston, K. J., Mechelli, A., Turner, R., & Price, C. J. (2000). Nonlinear responses in fmri: The balloon model, volterra kernels, and other hemodynamics. NeuroImage, vol. 12no. 4, 466–477.

Friston, K., Fletcher, P., Josephs, O., Holmes, A., Rugg, M., & Turner, R. (1998). Event-related fMRI: Characterizing differential responses. NeuroImage, vol. 7no. 1, 30–40.

Friston, K., Trujillo-Barreto, N., & Daunizeau, J. (2008). DEM: A variational treatment of dynamic systems. NeuroImage, vol. 41no. 3, 849–885.

Fu, Z., Tu, Y., Di, X., Biswal, B. B., Calhoun, V. D., & Zhang, Z. (2017). Associations between functional connectivity dynamics and BOLD dynamics are heterogeneous across brain networks. Frontiers in Human Neuroscience, vol. 11.

Gaudes, C. C., Karahanoglu, F. I., Lazeyras, F., & Ville, D. V. D. (2012). Structured sparse deconvolution for paradigm free mapping of functional

MRI data. 2012 9th IEEE International Symposium on Biomedical Imaging (ISBI).

Gaudes, C. C., Petridou, N., Dryden, I. L., Bai, L., Francis, S. T., & Gowland, P. A. (2010). Detection and characterization of single-trial fMRI bold responses: Paradigm free mapping. Human Brain Mapping, vol. 32no. 9, 1400–1418.

Gaudes, C. C., Petridou, N., Francis, S. T., Dryden, I. L., & Gowland, P. A. (2013). Paradigm free mapping with sparse regression automatically detects single-trial functional magnetic resonance imaging blood oxygenation level dependent responses. Human Brain Mapping, n/a–n/a.

Gaudes, C. C., Van De Ville, D., Petridou, N., Lazeyras, F., & Gowland, P. (2011). Paradigm-free mapping with morphological component analysis: Getting most out of fmri data. Wavelets and Sparsity XIV, vol. 8138, 81381K.

Gauthier, I., & Nelson, C. A. (2001). The development of face expertise. Current Opinion in Neurobiology, vol. 11no. 2, 219–224.

Geerligs, L., Gözükara, D., Oetringer, D., Campbell, K. L., van Gerven, M., & Güçlü, U. (2022). A partially nested cortical hierarchy of neural states underlies event segmentation in the human brain. eLife, vol. 11.

Geerligs, L., van Gerven, M., & Güçlü, U. (2021). Detecting neural state transitions underlying event segmentation. NeuroImage, vol. 236, 118085.

Gerchen, M. F., Bernal-Casas, D., & Kirsch, P. (2014). Analyzing task-dependent brain network changes by whole-brain psychophysiological interactions: A comparison to conventional analysis. Human Brain Mapping, vol. 35no. 10, 5071–5082.

Gitelman, D. R., Penny, W. D., Ashburner, J., & Friston, K. J. (2003). Modeling regional and psychophysiologic interactions in fMRI: The importance of hemodynamic deconvolution. NeuroImage, vol. 19no. 1, 200–207.

Glasser, M. F., Coalson, T. S., Bijsterbosch, J. D., Harrison, S. J., Harms, M. P., Anticevic, A., Essen, D. C. V., & Smith, S. M. (2018). Using temporal ICA to selectively remove global noise while preserving global signal in functional MRI data. NeuroImage, vol. 181, 692–717.

Glasser, M. F., Smith, S. M., Marcus, D. S., Andersson, J. L. R., Auerbach, E. J., Behrens, T. E. J., Coalson, T. S., Harms, M. P., Jenkinson, M., Moeller, S., Robinson, E. C., Sotiropoulos, S. N., Xu, J., Yacoub, E., Ugurbil, K., & Essen, D. C. V. (2016). The human connectome project's neuroimaging approach. Nature Neuroscience, vol. 19no. 9, 1175–1187.

Glerean, E., Salmi, J., Lahnakoski, J. M., Jääskeläinen, I. P., & Sams, M. (2012). Functional magnetic resonance imaging phase synchronization as a measure of dynamic functional connectivity. Brain Connectivity, vol. 2no. 2, 91–101.

Glover, G. H. (1999). Deconvolution of impulse response in event-related BOLD fMRI1. NeuroImage, vol. 9no. 4, 416–429.

Glover, G. H., Li, T.-Q., & Ress, D. (2000). Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. Magnetic Resonance in Medicine, vol. 44no. 1, 162–167.

Golestani, A. M., Chang, C., Kwinta, J. B., Khatamian, Y. B., & Chen, J. J. (2015). Mapping the end-tidal CO2 response function in the resting-state BOLD fMRI signal: Spatial specificity, test–retest reliability and effect of fMRI sampling rate. NeuroImage, vol. 104, 266–277.

Gonzalez-Castillo, J., Caballero-Gaudes, C., Topolski, N., Handwerker, D. A., Pereira, F., & Bandettini, P. A. (2019). Imaging the spontaneous flow of thought: Distinct periods of cognition contribute to dynamic functional connectivity during rest. NeuroImage, vol. 202, 116129.

Gonzalez-Castillo, J., Panwar, P., Buchanan, L. C., Caballero-Gaudes, C., Handwerker, D. A., Jangraw, D. C., Zachariou, V., Inati, S., Roopchansingh, V., Derbyshire, J. A., & Bandettini, P. A. (2016). Evaluation of multi-echo ICA denoising for task based fMRI studies: Block designs, rapid event-related designs, and cardiac-gated fMRI. NeuroImage, vol. 141, 452–468.

Gonzalez-Castillo, J., Saad, Z. S., Handwerker, D. A., Inati, S. J., Brenowitz, N., & Bandettini, P. A. (2012). Whole-brain, time-locked activation with simple tasks revealed using massive averaging and model-free analysis. Proceedings of the National Academy of Sciences of the United States of America, vol. 109, 5487–5492.

Goodyear, B. G., & Menon, R. S. (1998). Effect of luminance contrast on bold fmri response in human primary visual areas. Journal of Neurophysiology, vol. 79no. 4, 2204–2207.

Goutte, C., Nielsen, F. A., & Hansen, L. K. (2000). Modeling the haemodynamic response in fmri using smooth fir filters. IEEE transactions on medical imaging, vol. 19, 1188–1201.

Gramfort, A., Peyré, G., & Cuturi, M. (2015). Fast optimal transport averaging of neuroimaging data. In Lecture notes in computer science (pp. 261–272). Springer International Publishing.

Gramfort, A., Strohmeier, D., Haueisen, J., Hamalainen, M., & Kowalski, M. (2011). Functional brain imaging with m/EEG using structured sparsity in time-frequency dictionaries, 600–611.

Gratton, C., Dworetsky, A., Coalson, R. S., Adeyemo, B., Laumann, T. O., Wig, G. S., Kong, T. S., Gratton, G., Fabiani, M., Barch, D. M., Tranel, D., Miranda-Dominguez, O., Fair, D. A., Dosenbach, N. U., Snyder, A. Z., Perlmutter, J. S., Petersen, S. E., & Campbell, M. C. (2020). Removal of high frequency contamination from motion estimates in single-band fMRI saves data without biasing functional connectivity. NeuroImage, vol. 217, 116866.

Gratton, G., & Fabiani, M. (2001). The event-related optical signal: A new tool for studying brain function. International Journal of Psychophysiology, vol. 42no. 2, 109–121.

Gratton, G., Fabiani, M., Corballis, P. M., Hood, D. C., Goodman-Wood, M. R., Hirsch, J., Kim, K., Friedman, D., & Gratton, E. (1997). Fast and localized event-related optical signals (EROS) in the human occipital cortex: Comparisons with the visual evoked potential and fMRI. NeuroImage, vol. 6no. 3, 168–180.

Griffanti, L., Douaud, G., Bijsterbosch, J., Evangelisti, S., Alfaro-Almagro, F., Glasser, M. F., Duff, E. P., Fitzgibbon, S., Westphal, R., Carone, D., Beckmann, C. F., & Smith, S. M. (2017). Hand classification of fMRI ICA noise components. NeuroImage, vol. 154, 188–205.

Grill-Spector, K., & Malach, R. (2001). Fmr-adaptation: A tool for studying the functional properties of human cortical neurons. Acta Psychologica, vol. 107no. 1–3, 293–321.

Gueddari, L. E., Radhakrishna, C. G., Chouzenoux, E., & Ciuciu, P. (2021). Calibration-less multi-coil compressed sensing magnetic resonance image reconstruction based on OSCAR regularization. Journal of Imaging, vol. 7no. 3, 58.

Hammernik, K., Kustner, T., Yaman, B., Huang, Z., Rueckert, D., Knoll, F., & Akcakaya, M. (2023). Physics-driven deep learning for computational magnetic resonance imaging: Combining physics and machine learning for improved medical imaging. IEEE Signal Processing Magazine, vol. 40no. 1, 98–114.

Handwerker, D. A., Ollinger, J. M., & D'Esposito, M. (2004). Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. NeuroImage, vol. 21no. 4, 1639–1651.

Harrison, S. J., Bianchi, S., Heinzle, J., Stephan, K. E., Iglesias, S., & Kasper, L. (2021). A hilbert-based method for processing respiratory timeseries. NeuroImage, vol. 230, 117787.

Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., & Malach, R. (2004). Intersubject synchronization of cortical activity during natural vision. Science, vol. 303no. 5664, 1634–1640.

Hasson, U., Yang, E., Vallines, I., Heeger, D. J., & Rubin, N. (2008). A hierarchy of temporal receptive windows in human cortex. The Journal of Neuroscience, vol. 28no. 10, 2539–2550.

Havlicek, M., Friston, K. J., Jan, J., Brazdil, M., & Calhoun, V. D. (2011). Dynamic modeling of neuronal responses in fMRI using cubature kalman filtering. NeuroImage, vol. 56no. 4, 2109–2128.

Haxby, J. V., Hoffman, E. A., & Gobbini, M. (2000). The distributed human neural system for face perception. Trends in Cognitive Sciences, vol. 4no. 6, 223–233.

Henson, R., & Friston, K. (2007). Chapter 14 - convolution models for fmri. In K. Friston, J. Ashburner, S. Kiebel, T. Nichols, & W. Penny (Eds.), Statistical parametric mapping (pp. 178–192). Academic Press.

Hernandez-Garcia, L., & Ulfarsson, M. O. (2011). Neuronal event detection in fMRI time series using iterative deconvolution techniques. Magnetic Resonance Imaging, vol. 29no. 3, 353–364.

Honey, C. J., Thesen, T., Donner, T. H., Silbert, L. J., Carlson, C. E., Devinsky, O., Doyle, W. K., Rubin, N., Heeger, D. J., & Hasson, U. (2012). Slow cortical dynamics and the accumulation of information over long timescales. Neuron, vol. 76no. 2, 423–434.

Horien, C., Shen, X., Scheinost, D., & Constable, R. T. (2019). The individual functional connectome is unique and stable over months to years. NeuroImage, vol. 189, 676–687.

Hsu, W.-T., Rosenberg, M. D., Scheinost, D., Constable, R. T., & Chun, M. M. (2018). Resting-state functional connectivity predicts neuroticism and extraversion in novel individuals. Social Cognitive and Affective Neuroscience, vol. 13no. 2, 224–232.

Hu, X., Le, T. H., & U??urbil, K. (1997). Evaluation of the early response in fMRI in individual subjects using short stimulus duration. Magnetic Resonance in Medicine, vol. 37no. 6, 877–884.

Hutchison, R. M., Womelsdorf, T., Allen, E. A., Bandettini, P. A., Calhoun, V. D., Corbetta, M., Penna, S. D., Duyn, J. H., Glover, G. H., Gonzalez-Castillo, J., Handwerker, D. A., Keilholz, S., Kiviniemi, V., Leopold, D. A., de Pasquale, F., Sporns, O., Walter, M., & Chang, C. (2013). Dynamic functional connectivity: Promise, issues, and interpretations. NeuroImage, vol. 80, 360–378.

Hütel, M., Antonelli, M., Melbourne, A., & Ourselin, S. (2021). Hemodynamic matrix factorization for functional magnetic resonance imaging. NeuroImage, vol. 231, 117814.

Hütel, M., Melbourne, A., & Ourselin, S. (2018). Neural activation estimation in brain networks during task and rest using BOLD-fMRI. Medical Image Computing and Computer Assisted Intervention – MICCAI 2018, 215–222.

Hämäläinen, M., Hari, R., Ilmoniemi, R. J., Knuutila, J., & Lounasmaa, O. V. (1993). Magnetoencephalography—theory, instrumentation, and applications to noninvasive studies of the working human brain. Reviews of Modern Physics, vol. 65no. 2, 413–497.

Jafarpour, A., Buffalo, E. A., Knight, R. T., & Collins, A. G. (2022). Event segmentation reveals working memory forgetting rate. iScience, vol. 25no. 3, 103902.

Jain, S., Vo, V., Mahto, S., LeBel, A., Turek, J. S., & Huth, A. (2020). Interpretable multi-timescale models for predicting fmri responses to continuous natural speech. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, & H. Lin (Eds.), Advances in neural information processing systems (pp. 13738–13749, Vol. 33). Curran Associates, Inc.

Jalbrzikowski, M., Liu, F., Foran, W., Klei, L., Calabro, F. J., Roeder, K., Devlin, B., & Luna, B. (2020). Functional connectome fingerprinting accuracy in youths and adults is similar when examined on the same day and 1.5-years apart. Human Brain Mapping, vol. 41no. 15, 4187–4199.

Janati, H., Bazeille, T., Thirion, B., Cuturi, M., & Gramfort, A. (2020). Multi-subject MEG/EEG source imaging with sparse multi-task regression. NeuroImage, vol. 220, 116847.

Janati, H., Cuturi, M., & Gramfort, A. (2019, 16–18 Apr). Wasserstein regularization for sparse multi-task regression. In K. Chaudhuri & M. Sugiyama (Eds.), Proceedings of the twenty-second international

conference on artificial intelligence and statistics (pp. 1407–1416, Vol. 89). PMLR.

Janati, H., Cuturi, M., & Gramfort, A. (2022). Averaging spatio-temporal signals using optimal transport and soft alignments.

Jangraw, D. C., Finn, E. S., Bandettini, P. A., Landi, N., Sun, H., Hoeft, F., Chen, G., Pugh, K. R., & Molfese, P. J. (2023). Inter-subject correlation during long narratives reveals widespread neural correlates of reading ability. NeuroImage, vol. 282, 120390.

Jenkinson, M., Beckmann, C. F., Behrens, T. E., Woolrich, M. W., & Smith, S. M. (2012). FSL. NeuroImage, vol. 62no. 2, 782–790.

Jenkinson, M., & Smith, S. (2001). A global optimisation method for robust affine registration of brain images. Medical Image Analysis, vol. 5no. 2, 143–156.

Jewell, S. W., Hocking, T. D., Fearnhead, P., & Witten, D. M. (2019). Fast nonconvex deconvolution of calcium imaging data. Biostatistics, vol. 21no. 4, 709–726.

Jo, H. J., Saad, Z. S., Simmons, W. K., Milbury, L. A., & Cox, R. W. (2010). Mapping sources of correlation in resting state FMRI, with artifact detection and removal. NeuroImage, vol. 52no. 2, 571–582.

Jo, Y., Faskowitz, J., Esfahlani, F. Z., Sporns, O., & Betzel, R. F. (2021). Subject identification using edge-centric functional connectivity. NeuroImage, vol. 238, 118204.

Jorge, J., Figueiredo, P., van der Zwaag, W., & Marques, J. P. (2013). Signal fluctuations in fMRI data acquired with 2d-EPI and 3d-EPI at 7 tesla. Magnetic Resonance Imaging, vol. 31no. 2, 212–220.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. The Journal of Neuroscience, vol. 17no. 11, 4302–4311.

Kanwisher, N., & Yovel, G. (2006). The fusiform face area: A cortical region specialized for the perception of faces. Philosophical Transactions of the Royal Society B: Biological Sciences, vol. 361no. 1476, 2109–2128.

Karahanoglu, F. I., Grouiller, F., Gaudes, C. C., Seeck, M., Vulliemoz, S., & Van De Ville, D. (2013). Spatial mapping of interictal epileptic discharges in fMRI with total activation. 2013 IEEE 10th International Symposium on Biomedical Imaging, 1500–1503.

Karahanoglu, F. I., Bayram, İ., & Van De Ville, D. (2011). A signal processing approach to generalized 1-d total variation. IEEE Transactions on Signal Processing, vol. 59no. 11, 5265–5274.

Karahanoğlu, F. I., & Van De Ville, D. (2015). Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks. Nature Communications, vol. 6no. 1, 7751.

Karahanoğlu, F. I., Caballero-Gaudes, C., Lazeyras, F., & Van De Ville, D. (2013). Total activation: fMRI deconvolution through spatio-temporal regularization. NeuroImage, vol. 73, 121–134.

Karahanoğlu, F. I., & Van De Ville, D. (2017). Dynamics of large-scale fMRI networks: Deconstruct brain activity to build better models of brain function. Current Opinion in Biomedical Engineering, vol. 3, 28–36.

Kassinopoulos, M., & Mitsis, G. D. (2019). Identification of physiological response functions to correct for fluctuations in resting-state fMRI related to heart rate and respiration. NeuroImage, vol. 202, 116150.

Keilholz, S., Caballero-Gaudes, C., Bandettini, P., Deco, G., & Calhoun, V. (2017). Time-resolved resting-state functional magnetic resonance imaging analysis: Current status, challenges, and new directions. Brain Connectivity, vol. 7no. 8, 465–481.

Keilholz, S. D. (2014). The neural basis of time-varying resting-state functional connectivity. Brain Connectivity, vol. 4no. 10, 769–779.

Khalidov, I., Fadili, J., Lazeyras, F., Van De Ville, D., & Unser, M. (2011). Activelets: Wavelets for sparse representation of hemodynamic responses. Signal Processing, vol. 91no. 12, 2810–2821.

Kinany, N., Pirondini, E., Micera, S., & Van De Ville, D. (2020). Dynamic functional connectivity of resting-state spinal cord fMRI reveals fine-grained intrinsic architecture. Neuron, vol. 108no. 3, 424–435.e4.

Kleinschmidt, A., Obrig, H., Requardt, M., Merboldt, K.-D., Dirnagl, U., Villringer, A., & Frahm, J. (1996). Simultaneous recording of cerebral blood oxygenation changes during human brain activation by magnetic resonance imaging and near-infrared spectroscopy. Journal of Cerebral Blood Flow & Metabolism, vol. 16no. 5, 817–826.

Kowalski, M. (2009). Sparse regression using mixed norms. Applied and Computational Harmonic Analysis, vol. 27no. 3, 303–324.

Kruggel, F., & von Cramon, D. Y. (1999). Modeling the hemodynamic response in single-trial functional MRI experiments. Magnetic Resonance in Medicine, vol. 42no. 4, 787–797.

Kumar, M., Anderson, M., Antony, J., Baldassano, C., Brooks, P. P., Cai, M. B., Chen, P.-H. C., Ellis, C. T., Henselman-Petrusek, G., Huberdeau, D., Hutchinson, J. B., Li, Y. P., Lu, Q., Manning, J. R., Mennen, A. C., Nastase, S. A., Richard, H., Schapiro, A. C., Schuck, N. W., … Norman, K. A. (2020). BrainIAK: The brain imaging analysis kit.

Kundu, P., Brenowitz, N. D., Voon, V., Worbe, Y., Vértes, P. E., Inati, S. J., Saad, Z. S., Bandettini, P. A., & Bullmore, E. T. (2013). Integrated strategy for improving functional connectivity mapping using multiecho fMRI. Proceedings of the National Academy of Sciences, vol. 110no. 40, 16187–16192.

Kundu, P., Inati, S. J., Evans, J. W., Luh, W.-M., & Bandettini, P. A. (2012). Differentiating BOLD and non-BOLD signals in fMRI time series using multi-echo EPI. NeuroImage, vol. 60no. 3, 1759–1770.

Kundu, P., Voon, V., Balchandani, P., Lombardo, M. V., Poser, B. A., & Bandettini, P. A. (2017). Multi-echo fMRI: A review of applications in fMRI denoising and analysis of BOLD signals. NeuroImage, vol. 154, 59–80.

Kwong, K. K., Belliveau, J. W., Chesler, D. A., Goldberg, I. E., Weisskoff, R. M., Poncelet, B. P., Kennedy, D. N., Hoppel, B. E., Cohen, M. S., & Turner, R. (1992). Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. Proceedings of the National Academy of Sciences, vol. 89no. 12, 5675–5679.

Lake, E. M., Finn, E. S., Noble, S. M., Vanderwal, T., Shen, X., Rosenberg, M. D., Spann, M. N., Chun, M. M., Scheinost, D., & Constable, R. T. (2019). The functional brain organization of an individual allows prediction of measures of social abilities transdiagnostically in autism and attention-deficit/hyperactivity disorder. Biological Psychiatry, vol. 86no. 4, 315–326.

Lerner, Y., Honey, C. J., Silbert, L. J., & Hasson, U. (2011). Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. The Journal of Neuroscience, vol. 31no. 8, 2906–2915.

Liégeois, R., Li, J., Kong, R., Orban, C., Ville, D. V. D., Ge, T., Sabuncu, M. R., & Yeo, B. T. T. (2019). Resting brain dynamics at different timescales capture distinct aspects of human behavior. Nature Communications, vol. 10no. 1.

Liégeois, R., Yeo, B. T., & Ville, D. V. D. (2021). Interpreting null models of resting-state functional MRI dynamics: Not throwing the model out with the hypothesis. NeuroImage, vol. 243, 118518.

Liem, F., Varoquaux, G., Kynast, J., Beyer, F., Masouleh, S. K., Huntenburg, J. M., Lampe, L., Rahim, M., Abraham, A., Craddock, R. C., Riedel-Heller, S., Luck, T., Loeffler, M., Schroeter, M. L., Witte, A. V., Villringer, A., & Margulies, D. S. (2017). Predicting brain-age from multimodal imaging data captures cognitive impairment. NeuroImage, vol. 148, 179–188.

Lindquist, M. A., Xu, Y., Nebel, M. B., & Caffo, B. S. (2014). Evaluating dynamic bivariate correlations in resting-state fMRI: A comparison study and a new approach. NeuroImage, vol. 101, 531–546.

Liu, T. T. (2016). Noise contributions to the fMRI signal: An overview. NeuroImage, vol. 143, 141–151.

Liu, T. T., Li, B., Fernandez, B., & Banerjee, S. (2022). A geometric view of signal sensitivity metrics in multi-echo fmri. NeuroImage, vol. 259, 119409.

Liu, X., Chang, C., & Duyn, J. H. (2013). Decomposition of spontaneous brain activity into distinct fMRI co-activation patterns. Frontiers in Systems Neuroscience, vol. 7, 101.

Liu, X., & Duyn, J. H. (2013). Time-varying functional network information extracted from brief instances of spontaneous brain activity. Proceedings of the National Academy of Sciences, vol. 110no. 11, 4392–4397.

Liu, X., Zhang, N., Chang, C., & Duyn, J. H. (2018). Co-activation patterns in resting-state fMRI signals. NeuroImage, vol. 180, 485–494.

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. Nature, vol. 453no. 7197, 869–878.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. Nature, vol. 412no. 6843, 150–157.

Lombardo, M. V., Auyeung, B., Holt, R. J., Waldman, J., Ruigrok, A. N., Mooney, N., Bullmore, E. T., Baron-Cohen, S., & Kundu, P. (2016). Improving effect size estimation and statistical power with multi-echo fMRI and its impact on understanding the neural systems supporting mentalizing. NeuroImage, vol. 142, 55–66.

Lopes, R., Lina, J., Fahoum, F., & Gotman, J. (2012). Detection of epileptic activity in fMRI without recording the EEG. NeuroImage, vol. 60no. 3, 1867–1879.

López-Tapia, S., Molina, R., & Katsaggelos, A. K. (2021). Deep learning approaches to inverse problems in imaging: Past, present and future. Digital Signal Processing, vol. 119, 103285.

Loula, J., Varoquaux, G., & Thirion, B. (2018). Decoding fMRI activity in the time domain improves classification performance. NeuroImage, vol. 180, 203–210.

Lowe, M., Mock, B., & Sorenson, J. (1998). Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. NeuroImage, vol. 7no. 2, 119–132.

Lucas, A., Iliadis, M., Molina, R., & Katsaggelos, A. K. (2018). Using deep neural networks for inverse problems in imaging: Beyond analytical methods. IEEE Signal Processing Magazine, vol. 35no. 1, 20–36.

Lurie, D. J., Kessler, D., Bassett, D. S., Betzel, R. F., Breakspear, M., Kheilholz, S., Kucyi, A., Liégeois, R., Lindquist, M. A., McIntosh, A. R., Poldrack, R. A., Shine, J. M., Thompson, W. H., Bielczyk, N. Z., Douw, L., Kraft, D., Miller, R. L., Muthuraman, M., Pasquini, L., … Calhoun, V. D. (2020). Questions and controversies in the study of time-varying functional connectivity in resting fMRI. Network Neuroscience, vol. 4no. 1, 30–69.

Lynall, M.-E., Bassett, D. S., Kerwin, R., McKenna, P. J., Kitzbichler, M., Muller, U., & Bullmore, E. (2010). Functional connectivity and brain networks in schizophrenia. Journal of Neuroscience, vol. 30no. 28, 9477–9487.

Lynch, C. J., Power, J. D., Scult, M. A., Dubin, M., Gunning, F. M., & Liston, C. (2020). Rapid precision functional mapping of individuals using multi-echo fMRI. Cell Reports, vol. 33no. 12, 108540.

Madi, M. K., & Karameh, F. N. (2017). Hybrid cubature kalman filtering for identifying nonlinear models from sampled recording: Estimation of neuronal dynamics (X. Luo, Ed.). PLOS ONE, vol. 12no. 7, e0181513.

Majeed, W., Magnuson, M., Hasenkamp, W., Schwarb, H., Schumacher, E. H., Barsalou, L., & Keilholz, S. D. (2011). Spatiotemporal dynamics of low frequency BOLD fluctuations in rats and humans. NeuroImage, vol. 54no. 2, 1140–1150.

Malonek, D., & Grinvald, A. (1996). Interactions between electrical activity and cortical microcirculation revealed by imaging spectroscopy: Implications for functional brain mapping. Science, vol. 272no. 5261, 551–554.

Margulies, D. S., Ghosh, S. S., Goulas, A., Falkiewicz, M., Huntenburg, J. M., Langs, G., Bezgin, G., Eickhoff, S. B., Castellanos, F. X., Petrides, M., Jefferies, E., & Smallwood, J. (2016). Situating the default-mode network along a principal gradient of macroscale cortical organization. Proceedings of the National Academy of Sciences, vol. 113no. 44, 12574–12579.

Markello, R. D., Spreng, R. N., Luh, W.-M., Anderson, A. K., & Rosa, E. D. (2018). Segregation of the human basal forebrain using resting state functional MRI. NeuroImage, vol. 173, 287–297.

Marrelec, G., Benali, H., Ciuciu, P., & Poline, J.-B. (2002). Bayesian estimation of the hemodynamic response function in functional mri. AIP Conference Proceedings, vol. 617no. 1, 229–247.

Martino, A. D., Scheres, A., Margulies, D. S., Kelly, A. M. C., Uddin, L. Q., Shehzad, Z., Biswal, B., Walters, J. R., Castellanos, F. X., & Milham, M. P. (2008). Functional connectivity of human striatum: A resting state fMRI study. Cerebral Cortex, vol. 18no. 12, 2735–2747.

Mascali, D., Moraschi, M., DiNuzzo, M., Tommasin, S., Fratini, M., Gili, T., Wise, R. G., Mangia, S., Macaluso, E., & Giove, F. (2021). Evaluation of denoising strategies for task-based functional connectivity: Equalizing residual motion artifacts between rest and cognitively demanding tasks. Human Brain Mapping, vol. 42no. 6, 1805–1828.

McCann, M. T., Jin, K. H., & Unser, M. (2017). Convolutional neural networks for inverse problems in imaging: A review. IEEE Signal Processing Magazine, vol. 34no. 6, 85–95.

McGonigle, D., Howseman, A., Athwal, B., Friston, K., Frackowiak, R., & Holmes, A. (2000). Variability in fMRI: An examination of intersession differences. NeuroImage, vol. 11no. 6, 708–734.

Meinshausen, N., & Bühlmann, P. (2010). Stability selection. Journal of the Royal Statistical Society: Series B (Statistical Methodology), vol. 72no. 4, 417–473.

Menon, R. S., Luknowsky, D. C., & Gati, J. S. (1998). Mental chronometry using latency-resolved functional MRI. Proceedings of the National Academy of Sciences, vol. 95no. 18, 10902–10907.

Menon, R. S., Ogawa, S., Hu, X., Strupp, J. P., Anderson, P., & U??urbil, K. (1995). BOLD based functional MRI at 4 tesla includes a capillary bed contribution: Echo-planar imaging correlates with previous optical imaging using intrinsic signals. Magnetic Resonance in Medicine, vol. 33no. 3, 453–459.

Miezin, F., Maccotta, L., Ollinger, J., Petersen, S., & Buckner, R. (2000). Characterizing the hemodynamic response: Effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. NeuroImage, vol. 11no. 6, 735–759.

Miller, K. J., denNijs, M., Shenoy, P., Miller, J. W., Rao, R. P., & Ojemann, J. G. (2007). Real-time functional brain mapping using electrocorticography. NeuroImage, vol. 37no. 2, 504–507.

Miranda-Dominguez, O., Mills, B. D., Carpenter, S. D., Grant, K. A., Kroenke, C. D., Nigg, J. T., & Fair, D. A. (2014). Connectotyping: Model based fingerprinting of the functional connectome (S. Hayasaka, Ed.). PLoS ONE, vol. 9no. 11, e111048.

Moeller, S., Yacoub, E., Olman, C. A., Auerbach, E., Strupp, J., Harel, N., & Uğurbil, K. (2010). Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. Magnetic Resonance in Medicine, vol. 63no. 5, 1144–1153.

Moia, S., Termenon, M., Uruñuela, E., Chen, G., Stickland, R. C., Bright, M. G., & Caballero-Gaudes, C. (2021). ICA-based denoising strategies in breath-hold induced cerebrovascular reactivity mapping with multi echo BOLD fMRI. NeuroImage, vol. 233, 117914.

Monga, V., Li, Y., & Eldar, Y. C. (2021). Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. IEEE Signal Processing Magazine, vol. 38no. 2, 18–44.

Morris, P. G. (1999). Magnetic resonance imaging and magnetic resonance spectroscopy assessment of brain function in experimental animals and man. Journal of Psychopharmacology, vol. 13no. 4, 330–336.

Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., & Wang, X.-J. (2014). A hierarchy of intrinsic timescales across primate cortex. Nature Neuroscience, vol. 17no. 12, 1661–1663.

Muschelli, J., Nebel, M. B., Caffo, B. S., Barber, A. D., Pekar, J. J., & Mostofsky, S. H. (2014). Reduction of motion-related artifacts in resting state fMRI using aCompCor. NeuroImage, vol. 96, 22–35.

Nastase, S. A., Gazzola, V., Hasson, U., & Keysers, C. (2019). Measuring shared responses across subjects using intersubject correlation. Social Cognitive and Affective Neuroscience.

Neher, E., Sakmann, B., & Steinbach, J. H. (1978). The extracellular patch clamp: A method for resolving currents through individual open channels in biological membranes. Pflügers Archiv European Journal of Physiology, vol. 375no. 2, 219–228.

Nielsen, A. N., Greene, D. J., Gratton, C., Dosenbach, N. U. F., Petersen, S. E., & Schlaggar, B. L. (2018). Evaluating the prediction of brain maturity from functional connectivity after motion artifact denoising. Cerebral Cortex, vol. 29no. 6, 2455–2469.

Nir, Y., Mukamel, R., Dinstein, I., Privman, E., Harel, M., Fisch, L., Gelbard-Sagiv, H., Kipervasser, S., Andelman, F., Neufeld, M. Y., Kramer, U., Arieli, A., Fried, I., & Malach, R. (2008). Interhemispheric correlations of slow spontaneous neuronal fluctuations revealed in human sensory cortex. Nature Neuroscience, vol. 11no. 9, 1100–1108.

Nostro, A. D., Müller, V. I., Varikuti, D. P., Pläschke, R. N., Hoffstaedter, F., Langner, R., Patil, K. R., & Eickhoff, S. B. (2018). Predicting personality from network-based resting-state functional connectivity. Brain Structure and Function, vol. 223no. 6, 2699–2719.

Novelli, L., & Razi, A. (2022). A mathematical perspective on edge-centric brain functional connectivity. Nature Communications, vol. 13no. 1.

Nummenmaa, L., Glerean, E., Viinikainen, M., Jääskeläinen, I. P., Hari, R., & Sams, M. (2012). Emotions promote social interaction by synchronizing brain activity across individuals. Proceedings of the National Academy of Sciences, vol. 109no. 24, 9599–9604.

Ogawa, S., Lee, T. M., Kay, A. R., & Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Proceedings of the National Academy of Sciences, vol. 87no. 24, 9868–9872.

Ogawa, S., Tank, D. W., Menon, R., Ellermann, J. M., Kim, S. G., Merkle, H., & Ugurbil, K. (1992). Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. Proceedings of the National Academy of Sciences, vol. 89no. 13, 5951–5955.

Ongie, G., Jalal, A., Metzler, C. A., Baraniuk, R. G., Dimakis, A. G., & Willett, R. (2020). Deep learning techniques for inverse problems in imaging. IEEE Journal on Selected Areas in Information Theory, vol. 1no. 1, 39–56.

Ortelli, F., & van de Geer, S. (2019). Synthesis and analysis in total variation regularization.

Otazo, R., Candes, E., & Sodickson, D. K. (2015). Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components. Magnetic resonance in medicine, vol. 73no. 3, 1125–1136.

Pais-Roldán, P., Biswal, B., Scheffler, K., & Yu, X. (2018). Identifying respiration-related aliasing artifacts in the rodent resting-state fMRI. Frontiers in Neuroscience, vol. 12.

Parker, D. B., & Razlighi, Q. R. (2019). The benefit of slice timing correction in common fMRI preprocessing pipelines. Frontiers in Neuroscience, vol. 13.

Patriat, R., Molloy, E. K., & Birn, R. M. (2015). Using edge voxel information to improve motion regression for rs-fMRI connectivity studies. Brain Connectivity, vol. 5no. 9, 582–595.

Peña-Gómez, C., Avena-Koenigsberger, A., Sepulcre, J., & Sporns, O. (2017). Spatiotemporal network markers of individual variability in the human functional connectome. Cerebral Cortex, vol. 28no. 8, 2922–2934.

Penny, W., Ghahramani, Z., & Friston, K. (2005). Bilinear dynamical systems. Philosophical Transactions of the Royal Society B: Biological Sciences, vol. 360no. 1457, 983–993.

Perlbarg, V., Bellec, P., Anton, J.-L., Pélégrini-Issac, M., Doyon, J., & Benali, H. (2007). CORSICA: Correction of structured noise in fMRI by automatic identification of ICA components. Magnetic Resonance Imaging, vol. 25no. 1, 35–46.

Petridou, N., Gaudes, C. C., Dryden, I. L., Francis, S. T., & Gowland, P. A. (2013). Periods of rest in fMRI contain individual spontaneous events

which are related to slowly fluctuating spontaneous activity. Human Brain Mapping, vol. 34no. 6, 1319–1329.

Pidnebesna, A., Fajnerová, I., Horáček, J., & Hlinka, J. (2019). Estimating sparse neuronal signal from hemodynamic response: The mixture components inference approach. bioRxiv.

Pinto, J., Bright, M. G., Bulte, D. P., & Figueiredo, P. (2021). Cerebrovascular reactivity mapping without gas challenges: A methodological guide. Frontiers in Physiology, vol. 11.

Plitt, M., Barnes, K. A., Wallace, G. L., Kenworthy, L., & Martin, A. (2015). Resting-state functional connectivity predicts longitudinal change in autistic traits and adaptive functioning in autism. Proceedings of the National Academy of Sciences, vol. 112no. 48.

Polimeni, J. R., & Lewis, L. D. (2021). Imaging faster neural dynamics with fast fMRI: A need for updated models of the hemodynamic response. Progress in Neurobiology, vol. 207, 102174.

Poser, B. A., Versluis, M. J., Hoogduin, J. M., & Norris, D. G. (2006). BOLD contrast sensitivity enhancement and artifact reduction with multiecho EPI: Parallel-acquired inhomogeneity-desensitized fMRI. Magnetic Resonance in Medicine, vol. 55no. 6, 1227–1235.

Posse, S., Wiese, S., Gembris, D., Mathiak, K., Kessler, C., Grosse-Ruyken, M.-L., Elghahwagi, B., Richards, T., Dager, S. R., & Kiselev, V. G. (1999). Enhancement of BOLD-contrast sensitivity by single-shot multi-echo functional MR imaging. Magnetic Resonance in Medicine, vol. 42no. 1, 87–97.

Pourtois, G., Thut, G., Grave de Peralta, R., Michel, C., & Vuilleumier, P. (2005). Two electrophysiological stages of spatial orienting towards fearful faces: Early temporo-parietal activation preceding gain control in extrastriate visual cortex. NeuroImage, vol. 26no. 1, 149–163.

Power, J. D. (2017). A simple but useful way to assess fMRI scan qualities. NeuroImage, vol. 154, 150–158.

Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. NeuroImage, vol. 59no. 3, 2142–2154.

Power, J. D., Lynch, C. J., Silver, B. M., Dubin, M. J., Martin, A., & Jones, R. M. (2019). Distinctions among real and apparent respiratory motions in human fMRI data. NeuroImage, vol. 201, 116041.

Power, J. D., Plitt, M., Gotts, S. J., Kundu, P., Voon, V., Bandettini, P., & Martin, A. (2018). Ridding fMRI data of motion-related influences: Removal of signals with distinct spatial and physical bases in multiecho data. Proceedings of the National Academy of Sciences, vol. 115, E2105–E2114.

Power, J. D., Plitt, M., Laumann, T. O., & Martin, A. (2017). Sources and implications of whole-brain fMRI signals in humans. NeuroImage, vol. 146, 609–625.

Preti, M. G., Bolton, T. A., & Van De Ville, D. (2017). The dynamic functional connectome: State-of-the-art and perspectives. NeuroImage, vol. 160, 41–54.

Pruim, R. H., Mennes, M., Buitelaar, J. K., & Beckmann, C. F. (2015). Evaluation of ICA-AROMA and alternative strategies for motion artifact removal in resting state fMRI. NeuroImage, vol. 112, 278–287.

Pruim, R. H., Mennes, M., van Rooij, D., Llera, A., Buitelaar, J. K., & Beckmann, C. F. (2015). ICA-AROMA: A robust ICA-based strategy for removing motion artifacts from fMRI data. NeuroImage, vol. 112, 267–277.

Puckett, A. M., Bollmann, S., Poser, B. A., Palmer, J., Barth, M., & Cunnington, R. (2018). Using multi-echo simultaneous multi-slice (SMS) EPI to improve functional MRI of the subcortical nuclei of the basal ganglia at ultra-high field (7t). NeuroImage, vol. 172, 886–895.

Raccah, O., Doelling, K. B., Davachi, L., & Poeppel, D. (2022). Acoustic features drive event segmentation in speech. Journal of Experimental Psychology: Learning, Memory, and Cognition.

Raguet, H., Fadili, J., & Peyré, G. (2013). A generalized forward-backward splitting. SIAM Journal on Imaging Sciences, vol. 6no. 3, 1199–1226.

Raj, D., Anderson, A. W., & Gore, J. C. (2001). Respiratory effects in human functional magnetic resonance imaging due to bulk susceptibility changes. Physics in Medicine and Biology, vol. 46no. 12, 3331–3340.

Rajapakse, J. C., Kruggel, F., Maisog, J. M., & von Cramon, D. Y. (1998). Modeling hemodynamic response for analysis of functional MRI time-series. Human Brain Mapping, vol. 6no. 4, 283–300.

Rangaprakash, D., Wu, G.-R., Marinazzo, D., Hu, X., & Deshpande, G. (2018). Hemodynamic response function (hrf) variability confounds resting-state fmri functional connectivity. Magnetic Resonance in Medicine, vol. 80no. 4, 1697–1713.

Richter, W., Ugurbil, K., Georgopoulos, A., & Kim, S.-G. (1997). Time-resolved fMRI of mental rotation. NeuroReport, vol. 8no. 17, 3697–3702.

Riera, J. J., Watanabe, J., Kazuki, I., Naoki, M., Aubert, E., Ozaki, T., & Kawashima, R. (2004). A state-space model of the hemodynamic approach: Nonlinear filtering of BOLD signals. NeuroImage, vol. 21no. 2, 547–567.

Rolls, E. T., Cheng, W., & Feng, J. (2021). Brain dynamics: Synchronous peaks, functional connectivity, and its temporal variability. Human brain mapping, vol. 42, 2790–2801.

Rosenberg, M. D., Finn, E. S., Scheinost, D., Papademetris, X., Shen, X., Constable, R. T., & Chun, M. M. (2015). A neuromarker of sustained attention from whole-brain functional connectivity. Nature Neuroscience, vol. 19no. 1, 165–171.

Rosenberg, M. D., Scheinost, D., Greene, A. S., Avery, E. W., Kwon, Y. H., Finn, E. S., Ramani, R., Qiu, M., Constable, R. T., & Chun, M. M. (2020). Functional connectivity predicts changes in attention observed across minutes, days, and months. Proceedings of the National Academy of Sciences, vol. 117no. 7, 3797–3807.

Rossion, B., Schiltz, C., & Crommelinck, M. (2003). The functionally defined right occipital and fusiform "face areas" discriminate novel from visually familiar faces. NeuroImage, vol. 19no. 3, 877–883.

Rotshtein, P., Henson, R. N. A., Treves, A., Driver, J., & Dolan, R. J. (2004). Morphing marilyn into maggie dissociates physical and identity face representations in the brain. Nature Neuroscience, vol. 8no. 1, 107–113.

Ruiz-Euler, H.-C., Marques, J. R. F., & Kappen, H. J. (2018). Nonlinear deconvolution by sampling biophysically plausible hemodynamic models.

Ryali, S., Chen, T., Supekar, K., & Menon, V. (2012). Estimation of functional connectivity in fmri data using stability selection-based sparse partial correlation with elastic net penalty. NeuroImage, vol. 59no. 4, 3852–3861.

Sadaghiani, S., Uğurbil, K., & Uludağ, K. (2009). Neural activity-induced modulation of BOLD poststimulus undershoot independent of the positive signal. Magnetic Resonance Imaging, vol. 27no. 8, 1030–1038.

Salimi-Khorshidi, G., Douaud, G., Beckmann, C. F., Glasser, M. F., Griffanti, L., & Smith, S. M. (2014). Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers. NeuroImage, vol. 90, 449–468.

Salmi, J., Metwaly, M., Tohka, J., Alho, K., Leppämäki, S., Tani, P., Koski, A., Vanderwal, T., & Laine, M. (2020). ADHD desynchronizes brain activity during watching a distracted multi-talker conversation. NeuroImage, vol. 216, 116352.

Salvador, R., Suckling, J., Coleman, M. R., Pickard, J. D., Menon, D., & Bullmore, E. (2005). Neurophysiological architecture of functional magnetic resonance images of human brain. Cerebral Cortex, vol. 15no. 9, 1332–1342.

Samara, A., Eilbott, J., Margulies, D. S., Xu, T., & Vanderwal, T. (2023). Cortical gradients during naturalistic processing are hierarchical and modality-specific. NeuroImage, vol. 271, 120023.

Sava-Segal, C., Richards, C., Leung, M., & Finn, E. S. (2022). Individual variability in neural event segmentation reflects stimulus content and interpretation.

Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A. J., Eickhoff, S. B., & Yeo, B. T. T. (2017). Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. Cerebral Cortex, vol. 28no. 9, 3095–3114.

Schwarz, G. (1978). Estimating the dimension of a model. The Annals of Statistics, vol. 6no. 2, 461–464.

Setsompop, K., Gagoski, B. A., Polimeni, J. R., Witzel, T., Wedeen, V. J., & Wald, L. L. (2011). Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g -factor penalty. Magnetic Resonance in Medicine, vol. 67no. 5, 1210–1224.

Shmuel, A., Augath, M., Oeltermann, A., & Logothetis, N. K. (2006). Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area v1. Nature Neuroscience, vol. 9no. 4, 569–577.

Shmueli, K., van Gelderen, P., de Zwart, J. A., Horovitz, S. G., Fukunaga, M., Jansma, J. M., & Duyn, J. H. (2007). Low-frequency fluctuations in the cardiac rate as a source of variance in the resting-state fMRI BOLD signal. NeuroImage, vol. 38no. 2, 306–320.

Silva, M., Baldassano, C., & Fuentemilla, L. (2019). Rapid memory reactivation at movie event boundaries promotes episodic encoding. The Journal of Neuroscience, vol. 39no. 43, 8538–8548.

Smith, S. M., Beckmann, C. F., Ramnani, N., Woolrich, M. W., Bannister, P. R., Jenkinson, M., Matthews, P. M., & McGonigle, D. J. (2005). Variability in fMRI: A re-examination of inter-session differences. Human Brain Mapping, vol. 24no. 3, 248–257.

Smith, S. M., Miller, K. L., Moeller, S., Xu, J., Auerbach, E. J., Woolrich, M. W., Beckmann, C. F., Jenkinson, M., Andersson, J., Glasser, M. F., Essen, D. C. V., Feinberg, D. A., Yacoub, E. S., & Ugurbil, K. (2012). Temporally-independent functional modes of spontaneous brain activity. Proceedings of the National Academy of Sciences, vol. 109no. 8, 3131–3136.

Smyser, C. D., Inder, T. E., Shimony, J. S., Hill, J. E., Degnan, A. J., Snyder, A. Z., & Neil, J. J. (2010). Longitudinal analysis of neural network development in preterm infants. Cerebral Cortex, vol. 20no. 12, 2852–2862.

Smyser, C. D., Snyder, A. Z., & Neil, J. J. (2011). Functional connectivity mri in infants: Exploration of the functional organization of the developing brain. NeuroImage, vol. 56no. 3, 1437–1452.

Speer, N. K., Reynolds, J. R., Swallow, K. M., & Zacks, J. M. (2009). Reading stories activates neural representations of visual and motor experiences. Psychological Science, vol. 20no. 8, 989–999.

Sporns, O., Faskowitz, J., Teixeira, A. S., Cutts, S. A., & Betzel, R. F. (2021). Dynamic expression of brain functional systems disclosed by fine-scale analysis of edge time series. Network neuroscience (Cambridge, Mass.), vol. 5, 405–433.

Sreenivasan, K. R., Havlicek, M., & Deshpande, G. (2015). Nonparametric hemodynamic deconvolution of fMRI using homomorphic filtering. IEEE Transactions on Medical Imaging, vol. 34no. 5, 1155–1163.

Sripada, C., Rutherford, S., Angstadt, M., Thompson, W. K., Luciana, M., Weigard, A., Hyde, L. H., & Heitzeg, M. (2019). Prediction of neurocognition in youth from resting state fMRI. Molecular Psychiatry, vol. 25no. 12, 3413–3421.

St-Yves, G., Allen, E. J., Wu, Y., Kay, K., & Naselaris, T. (2023). Brain-optimized deep neural network models of human visual areas learn non-hierarchical representations. Nature Communications, vol. 14no. 1.

Svaldi, D. O., Goñi, J., Abbas, K., Amico, E., Clark, D. G., Muralidharan, C., Dzemidzic, M., West, J. D., Risacher, S. L., Saykin, A. J., & Apostolova, L. G. (2021). Optimizing differential identifiability improves connectome predictive modeling of cognitive deficits from functional connectivity in alzheimer's disease. Human Brain Mapping, vol. 42no. 11, 3500–3516.

Saad, Z. S., Ropella, K. M., Cox, R. W., & DeYoe, E. A. (2001). Analysis and use of FMRI response delays. Human Brain Mapping, vol. 13no. 2, 74–93.

Tagliazucchi, E., Balenzuela, P., Fraiman, D., & Chialvo, D. R. (2012). Criticality in large-scale brain fmri dynamics unveiled by a novel point process analysis. Frontiers in physiology, vol. 3, 15.

Tagliazucchi, E., Balenzuela, P., Fraiman, D., Montoya, P., & Chialvo, D. R. (2011). Spontaneous bold event triggered averages for estimating functional connectivity at resting state. Neuroscience letters, vol. 488, 158–163.

Tagliazucchi, E., Siniatchkin, M., Laufs, H., & Chialvo, D. R. (2016). The voxel-wise functional connectome can be efficiently derived from co-activations in a sparse spatio-temporal point-process. Frontiers in Neuroscience, vol. 10, 381.

Tan, F. M., Caballero-Gaudes, C., Mullinger, K. J., Cho, S.-Y., Zhang, Y., Dryden, I. L., Francis, S. T., & Gowland, P. A. (2017). Decoding fMRI events in sensorimotor motor network using sparse paradigm free mapping and activation likelihood estimates. Human Brain Mapping, vol. 38no. 11, 5778–5794.

Tarun, A., Wainstein-Andriano, D., Sterpenich, V., Bayer, L., Perogamvros, L., Solms, M., Axmacher, N., Schwartz, S., & Ville, D. V. D. (2021). NREM sleep stages specifically alter dynamical integration of large-scale brain networks. iScience, vol. 24no. 1, 101923.

The Tedana Community, Ahmed, Z., Bandettini, P. A., Bottenhorn, K. L., Caballero-Gaudes, C., Dowdle, L. T., DuPre, E., Gonzalez-Castillo, J., Handwerker, D., Heunis, S., Kundu, P., Laird, A. R., Markello, R., Markiewicz, C. J., Maullin-Sapey, T., Moia, S., Salo, T., Staden, I., Teves, J., ... Whitaker, K. (2022). Me-ica/tedana: 0.0.12.

Thompson, G. J., Merritt, M. D., Pan, W.-J., Magnuson, M. E., Grooms, J. K., Jaeger, D., & Keilholz, S. D. (2013). Neural correlates of time-varying functional connectivity in the rat. NeuroImage, vol. 83, 826–836.

Thompson, G. J., Pan, W.-J., Magnuson, M. E., Jaeger, D., & Keilholz, S. D. (2014). Quasi-periodic patterns (QPP): Large-scale dynamics in resting state fMRI that correlate with local infraslow electrical activity. NeuroImage, vol. 84, 1018–1031.

Thompson, W. H., & Fransson, P. (2015). The mean–variance relationship reveals two possible strategies for dynamic brain connectivity analysis in fMRI. Frontiers in Human Neuroscience, vol. 9.

Tian, Y., Margulies, D. S., Breakspear, M., & Zalesky, A. (2020). Topographic organization of the human subcortex unveiled with functional connectivity gradients. Nature Neuroscience, vol. 23no. 11, 1421–1432.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society: Series B (Methodological), vol. 58no. 1, 267–288.

Triantafyllou, C., Hoge, R., Krueger, G., Wiggins, C., Potthast, A., Wiggins, G., & Wald, L. (2005). Comparison of physiological noise at 1.5 t, 3 t

and 7 t and optimization of fMRI acquisition parameters. NeuroImage, vol. 26no. 1, 243–250.

Turner, R., Bihan, D. L., Moonen, C. T. W., Despres, D., & Frank, J. (1991). Echo-planar time course MRI of cat brain oxygenation changes. Magnetic Resonance in Medicine, vol. 22no. 1, 159–166.

Uruñuela, E. (2023). Me-ica/rica: V1.0.17.

Uruñuela, E., Bolton, T. A., Van De Ville, D., & Caballero-Gaudes, C. (2023). Hemodynamic deconvolution demystified: Sparsity-driven regularization at work. Aperture Neuro, vol. 3, 1–25.

Uruñuela, E., Gonzalez-Castillo, J., Zheng, C., Bandettini, P., & Caballero-Gaudes, C. (2022). Whole-brain multivariate hemodynamic deconvolution for multi-echo fMRI with stability selection.

Uruñuela, E., Jones, S., Crawford, A., Shin, W., Oh, S., Lowe, M., & Caballero-Gaudes, C. (2020). Stability-based sparse paradigm free mapping algorithm for deconvolution of functional MRI data. 2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC).

Uruñuela, E., Moia, S., & Caballero-Gaudes, C. (2021). A low rank and sparse paradigm free mapping algorithm for deconvolution of fMRI data. 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI).

Uruñuela-Tremiño, E., Moia, S., Gonzalez-Castillo, J., & Caballero-Gaudes, C. (2019). Deconvolution of multi-echo functional MRI data with multivariate multi-echo sparse paradigm free mapping. ISMRM 27th Annual Meeting and Exhibition.

van der Zwaag, W., Francis, S., Head, K., Peters, A., Gowland, P., Morris, P., & Bowtell, R. (2009). fMRI at 1.5, 3 and 7 t: Characterising BOLD signal changes. NeuroImage, vol. 47no. 4, 1425–1434.

Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., & Ugurbil, K. (2013). The WU-minn human connectome project: An overview. NeuroImage, vol. 80, 62–79.

Vanderwal, T., Eilbott, J., Finn, E. S., Craddock, R. C., Turnbull, A., & Castellanos, F. X. (2017). Individual differences in functional connectivity during naturalistic viewing conditions. NeuroImage, vol. 157, 521–530.

van Oort, E. S. B., Mennes, M., Navarro Schröder, T., Kumar, V. J., Zaragoza Jimenez, N. I., Grodd, W., Doeller, C. F., & Beckmann, C. F. (2018). Functional parcellation using time courses of instantaneous connectivity. NeuroImage, vol. 170, 31–40.

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, vol. 154no. 1-2, 101–104.

Waller, L., Walter, H., Kruschwitz, J. D., Reuter, L., Müller, S., Erk, S., & Veer, I. M. (2017). Evaluating the replicability, specificity, and generalizability of connectome fingerprints. NeuroImage, vol. 158, 371–377.

Wang, F., Eljarrat, A., Müller, J., Henninen, T. R., Erni, R., & Koch, C. T. (2020). Multi-resolution convolutional neural networks for inverse problems. Scientific Reports, vol. 10no. 1.

Wang, J., Ren, Y., Hu, X., Nguyen, V. T., Guo, L., Han, J., & Guo, C. C. (2017). Test-retest reliability of functional connectivity networks during naturalistic fMRI paradigms. Human Brain Mapping, vol. 38no. 4, 2226–2241.

White, B. R., & Culver, J. P. (2010). Quantitative evaluation of high-density diffuse optical tomography: In vivo resolution and mapping performance. Journal of Biomedical Optics, vol. 15no. 02, 1.

Wilson, S. M., Molnar-Szakacs, I., & Iacoboni, M. (2007). Beyond superior temporal cortex: Intersubject correlations in narrative speech comprehension. Cerebral Cortex, vol. 18no. 1, 230–242.

Wise, R. G., Ide, K., Poulin, M. J., & Tracey, I. (2004). Resting fluctuations in arterial carbon dioxide induce significant low frequency variations in BOLD signal. NeuroImage, vol. 21no. 4, 1652–1664.

Yaesoubi, M., Allen, E. A., Miller, R. L., & Calhoun, V. D. (2015). Dynamic coherence analysis of resting fMRI data to jointly capture state-based phase, frequency, and time-domain information. NeuroImage, vol. 120, 133–142.

Yamashita, M., Yoshihara, Y., Hashimoto, R., Yahata, N., Ichikawa, N., Sakai, Y., Yamada, T., Matsukawa, N., Okada, G., Tanaka, S. C., Kasai, K., Kato, N., Okamoto, Y., Seymour, B., Takahashi, H., Kawato, M., & Imamizu, H. (2018). A prediction model of working memory across health and psychiatric disease using whole-brain functional connectivity. eLife, vol. 7.

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, vol. 106no. 3, 1125–1165.

Yeşilyurt, B., Uğurbil, K., & Uludağ, K. (2008). Dynamics and nonlinearities of the BOLD response at very short stimulus durations. Magnetic Resonance Imaging, vol. 26no. 7, 853–862.

Yousefi, B., & Keilholz, S. (2021). Propagating patterns of intrinsic activity along macroscale gradients coordinate functional connections across the whole brain. NeuroImage, vol. 231, 117827.

Yousefi, B., Shin, J., Schumacher, E. H., & Keilholz, S. D. (2018). Quasi-periodic patterns of intrinsic brain activity in individuals and their relationship to global signal. NeuroImage, vol. 167, 297–308.

Yovel, G., & Kanwisher, N. (2004). Face perception. Neuron, vol. 44no. 5, 889–898.

Yuan, R., Di, X., Taylor, P. A., Gohel, S., Tsai, Y.-H., & Biswal, B. B. (2015). Functional topography of the thalamocortical system in human. Brain Structure and Function, vol. 221no. 4, 1971–1984.

Zacks, J. M., Kumar, S., Abrams, R. A., & Mehta, R. (2009). Using movement and intentions to understand human activity. Cognition, vol. 112no. 2, 201–216.

Zacks, J. M., Speer, N. K., Vettel, J. M., & Jacoby, L. L. (2006). Event understanding and memory in healthy aging and dementia of the alzheimer type. Psychology and Aging, vol. 21no. 3, 466–482.

Zalla, T., Verlut, I., Franck, N., Puzenat, D., & Sirigu, A. (2004). Perception of dynamic action in patients with schizophrenia. Psychiatry Research, vol. 128no. 1, 39–51.

Zhang, C., Baum, S. A., Adduru, V. R., Biswal, B. B., & Michael, A. M. (2018). Test-retest reliability of dynamic functional connectivity in resting state fMRI. NeuroImage, vol. 183, 907–918.

Zhang, X., Pan, W.-J., & Keilholz, S. D. (2020). The relationship between BOLD and neural activity arises from temporally sparse events. NeuroImage, vol. 207, 116390.

Zhen, Z., Fang, H., & Liu, J. (2013). The hierarchical brain network for face recognition (M. Ptito, Ed.). PLoS ONE, vol. 8no. 3, e59886.

Zhou, T., & Tao, D. (2011). Godec: Randomized low-rank & sparse matrix decomposition in noisy case. Proceedings of the 28th International Conference on Machine Learning, ICML 2011.

Zibetti, M. V. W., Helou, E. S., Regatte, R. R., & Herman, G. T. (2018). Monotone FISTA with variable acceleration for compressed sensing magnetic resonance imaging. IEEE transactions on computational imaging, vol. 5no. 1, 109–119.

Zoeller, D., Sandini, C., Karahanoglu, F. I., Schaer, M., Eliez, S., & Van De Ville, D. (2019). Large-scale brain network dynamics provide a measure of psychosis and anxiety in 22q11.2 deletion syndrome. Biological Psychiatry: Cognitive Neuroscience and Neuroimaging, vol. 4, 881–892.

Zoller, D. M., Bolton, T. A. W., Karahanoglu, F. I., Eliez, S., Schaer, M., & Ville, D. V. D. (2019). Robust recovery of temporal overlap between network activity using transient-informed spatio-temporal regression. IEEE Transactions on Medical Imaging, vol. 38no. 1, 291–302.

Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. Journal of the Royal Statistical Society: Series B (Statistical Methodology), vol. 67no. 2, 301–320.

Zvolanek, K. M., Moia, S., Dean, J. N., Stickland, R. C., Caballero-Gaudes, C., & Bright, M. G. (2023). Comparing end-tidal CO2, respiration volume per time (RVT), and average gray matter signal for mapping cerebrovascular reactivity amplitude and delay with breath-hold task BOLD fMRI. NeuroImage, vol. 272, 120038.