

From Compressed Sensing to deep learning based methods for functional Magnetic Resonance Imaging reconstruction at ultra-high field

De l'échantillonnage compressé à l'apprentissage profond pour la reconstruction d'Imagerie par Résonance Magnétique fonctionelle à très haut champ

Thèse de doctorat de l'université Paris-Saclay

École doctorale n° 575, Electrical, Optical, Bio : physics and Engineering (EOBE) Spécialité de doctorat: Physique et imagerie médicale Graduate School : Sciences de l'ingénierie et des systèmes Référent : Faculté des sciences d'Orsay

Thèse préparée dans l'unité de recherche **NeuroSpin** (Université Paris-Saclay, CEA), sous la direction de **Philippe Ciuciu**, Directeur de recherche et le co-encadrement de **Alexandre Vignaud**, Directeur de recherche

Thèse soutenue à Paris-Saclay, le 24 mars 2025, par

Pierre-Antoine COMBY

Composition du jury

Membres du jury avec voix délibérative

Emilie Chouzenoux	Présidente
Saclay	
Mehmet Akçakaya	Rapporteur & Examinateur
Associate Professor, University of Minnesota	
Mark Chiew	Rapporteur & Examinateur
Associate Professor, University of Toronto	
Florent Meyniel	Examinateur
Directeur de recherche, Inserm/CEA, Université Paris-Saclay	

NNT: 2025UPAST031

自他共深

"Mutual help and prosperity" — Jigorō Kanō

@ Pierre-Antoine Comby, 2025

Remerciements

Les travaux de recherche présentés dans ce manuscrit sont le fruit de trois ans et demi de travail, et ne sauraient se réduire uniquement à ce dernier. Plus que l'écriture de code LaTeX et Python, la création de belles figures et logos, ce que je voudrais mettre en avant ici sont les rencontres, les échanges, les discussions, les débats, les rires et pas les larmes qui ont jalonné ce parcours.

Mes premiers remerciements vont à mes directeurs de thèse, Philippe Ciuciu et Alexandre Vignaud. Philippe, Alexandre, c'est grâce à vous que j'ai pu découvrir le monde de la recherche, ses exigences, et ses avantages. Merci pour votre confiance, votre soutien et vos conseils, et surtout, surtout, pour votre patience.

Avec le temps, j'ai eu aussi la chance de bénéficier de l'expertise des membres permanents METRIC à Neurospin (Alexandre, Nicolas, Vincent, Frank, Alexis, Caroline, Aurélien) et MIND à Neurospin et à l'Inria (Philippe, Alexandre, Thomas, Bertrand, Demian). Merci pour vos questions et remarques sur mes travaux, mais aussi pour les discussions sur tout et rien le midi et au café, merci pour les retraites, les after-works et les karaokés. Il n'y a pas de bonne recherche sans une bonne ambiance.

Je voudrais aussi remercier tous les doctorants et postdoctorants et stagiaires avec qui j'ai partagé quelques mois ou années : Zaccharie, Guillaume, Redouane, Zaineb, Chaithya, Matthieu, Benoit, Julia, Antoine, Raphael, Alexis, Matthieu, vous m'avez montré ce qui m'attendait et aider à traverser toutes les étapes de la thèse. Joseph, Joseph, Anna, Emna, Natalia, Soukaina, Caini, Asma, vous avez été mes compagnons de route, et je vous souhaite tout le meilleur pour la votre. J'ai également eu la chance d'encadrer Benjamin et Marie pour leur stage, cette première expérience "de l'autre côté du bureau" m'a grandement appris et c'est aussi grâce à vous que j'ai pu explorer tant de sujet différents.

Et puis, il y a tous ceux qui ont fait que ces trois années ne se résument pas qu'à une chaise et un bureau. Je pense aux personnes qui ont partagé avec moi des moments de vie, et qui ont rendu ces années plus belles. Merci à mes amis, Kchanais, scouts ou judoka, pour tous les moments passés ensemble, où l'on n'a plus à penser, mais juste à profiter.

Enfin, je voudrais remercier ma famille : merci à mes parents et grands-parents pour m'avoir transmis les valeurs et l'éducation qui m'ont permis d'arriver ici. Merci à mon frère et à mes deux sœurs pour votre soutien et votre amour. Merci aussi à la famille de Marine pour m'avoir accueilli et soutenu pendant toutes ces années. Surtout, je te remercie, Marine, pour ta patience et ton soutien inconditionnel. C'est grâce à toi que j'ai pu me concentrer sur ma thèse, et c'est grâce à toi que je suis devenu la personne que je suis aujourd'hui. Merci pour tout ce que tu fais pour moi, et pour tout ce que tu es.

Contents

Re	emero	ciements	v
Re	ésumo Mot Con Ape	é en Français ivation et Contexte	1 1 2 4
G	enera Con Con The	I Introduction text and Motivation	5 5 6 8
N	otatio	on and a second s	11
I	Int	troduction to functional MRI	13
1	Fun 1.1 1.2	ctional MRI: Spying on the brain with magnets Source of fMRI Signal: The BOLD contrast	 15 16 17 17 17 19 19 19 20
	1.3	1.2.2 From Bloch equations to the Fourier acquisition model Capturing the signal Image Formation and k-space encoding Image Formation and k-space encoding Image Formation 1.2.3 Pulse Sequences for fMRI Echo mechanism and Bold sensitivity Image Formation Constraints on the pulse sequence Image Formation Gradient Recalled Echo Imaging FMRI in practice Image Formation	 21 21 22 24 24 25 26 26 27
		1.3.1 the fMRI experiment fMRI equipment fmRI software pipeline	27 28 28

Contents

		1.3.2 Noise and artifacts sources and their mitigation strategies 29 Thermal Naise and SNP
		Dherricherical Meire
		Physiological Noise
		Field Inhomogeneities
		1.3.3 fMRI Analysis with the General Linear Model
	1.4	Conclusion
2	Gett	ting the fMRI signal faster 35
	2.1	Parallel imaging
		2.1.1 GRAPPA 36
		SNR in Parallel Imaging
		2.1.2 SENSE
		2.1.3 2D vs 3D imaging
	2.2	Non Cartesian imaging
		2.2.1 Spiral trajectories for fMRI 38
		2.2.2 Radial-Like trajectories for fMRI
	2.3	Compressed Sensing
		2.3.1 Sparsity
		2.3.2 Variable Density Sampling
		2.3.3 Optimizing the sampling trajectories
		2.3.4 Compressed sensing in practice
		Solvable formulation
		Analysis vs Synthesis 44
		2.3.5 Combining CS and PI 44
	2.4	Accelerating fMRI data acquisition using non-Cartesian compressed sensing 44
	$\frac{2.1}{2.5}$	Conclusion 46
	2.0	Concrusion

II Reconstructing functional MRI images

3	Ima	ge reco	onstruction methods for fMRI	49
	3.1	Comp	ressed Sensing and fMRI	49
		3.1.1	General problem formulation	49
		3.1.2	Frame-wise reconstruction	50
			Warm start strategies	51
		3.1.3	Time-aware reconstruction	51
			Low-rank approaches	51
			Temporal Fourier sparsity.	52
			Constrained Formulation.	52
			Penalized Formulation.	52
			Equivalence between the two formulations.	52
	3.2	Furthe	er extensions	53
		3.2.1	Other explicit priors	53
			Explicit rank prior (kt-FASTER).	53
			Matrix Factorization.	53
			Attachment to low resolution.	54
			Temporal subspace smoothness.	54

47

	3.3	Conclusion	54
4	Dee	ep Learning for MRI and functional MRI	57
	4.1	Deep Learning for inverse problems	57
		4.1.1 Supervised Learning	58
		4.1.2 Unsupervised and self-supervised learning	58
	4.2	Neural Network architecture for inverse problems	59
		4.2.1 U-NET Network as a building block	59
		4.2.2 Algorithm unrolling	60
		4.2.3 Plug-and-Play approaches	61
	4.3	Application to MRI And functional MRI	62
		4.3.1 Challenges for Deep Learning in MRI	62
		4.3.2 Deep Learning methods for fMRI reconstruction is even more chal-	
		lenging	63
		4.3.3 An hybrid solution with Plug-and-Play priors	63
	4.4	Conclusion	64
Su	ımma	ary of concerns	65

III Methodological developments

67

5	MR	I-NUFF	T: Doing non-Cartesian MRI has never been easier.	69
	Cont	text and	d summary	70
	5.1	Featur	es	70
		5.1.1	NUFFT Library compatibility	70
		5.1.2	Extended Fourier Model	71
			Parallel Imaging Model	72
			Off-resonance correction model	72
		5.1.3	Density Compensation	73
		5.1.4	Trajectories generation and expansions	74
		5.1.5	Auto-differentiation for data and sampling pattern	74
	5.2	Bench	mark of existing NUFFT libraries	74
	5.3	Examp	ole of Non-Cartesian, Non-Fourier Reconstruction using MRI-NUFFT	
		and co	mpressed sensing	74
		5.3.1	Material and methods	74
		5.3.2	Results	76
	5.4	Conclu	nsion	76
6	Ben	chmark	king local low-rank denoising methods for fMRI	79
	Cont	text and	d Summary	79
	6.1	Introd	uction	80
	6.2	Denois	sing methods using patch based PCA	81
		6.2.1	From global to local low rank assumption.	81
		6.2.2	Local low rank formalism	81
		6.2.3	Comparison of local low-rank methods	82

	6.3	Materi	ial and Methods	83
		6.3.1	Acquisition	83
		6.3.2	Preprocessing	84
		6.3.3	Statistical Analysis	84
	6.4	Result	S	84
		6.4.1	Single subject analysis	84
		6.4.2	"Group-level" analysis	86
		6.4.3	Noise map estimation	87
	6.5	Conclu	ision	87
7	No	Ground	-Truth ? Let's build some	91
	Cont	text and	d Summary	92
	7.1	Motiva	ation	93
	7.2	Existin	ng software and their limitations	94
	7.3	Extend	ded Fourier model for fMRI acquisition	96
		7.3.1	Single shot acquisition	96
		7.3.2	BOLD as a TE-sensitive change between shots	98
		7.3.3	Noise and SNR calibration	98
			Extending the signal model	99
		7.3.4	Summary of general hypotheses on the acquisition model	100
	7.4	Main o	characteristics of SNAKE implementation	100
		7.4.1	Modular Approach	100
		7.4.2	Performance, reproducibility and scalability made easier for neuro-	
			scientists	102
			Performance	102
			Reproducibility	102
			Scalability and interoperatibility	103
	7.5	Numer	rical experiments	103
		7.5.1	Acquisition scenarios	103
			Working hypotheses to simplify acquisition	103
			Scenario S1: 3D fully sampled Cartesian readout	104
			Scenario S2: 3D under-sampled stack of spirals (SoS) readout with	
			VDS acceleration along the stacking axis	105
			Scenario S3: Fully 3D SPARKLING under-sampled readout	105
		7.5.2	Reconstruction Strategies for Scenario 2 & 3	106
		7.5.3	Evaluation methods	107
	7.6	Result	8	107
		7.6.1	Scenario S1	107
		7.6.2	Scenario S2	110
			T_2^* relaxation must be taken into account for the long readout in	
			non-Cartesian acquisitions	111
			The refined initialization is key in dynamic acquisition for improved	
			performances	112
			Image quality is not a right proxy for good statistical performances	112
		7.6.3	Scenario S3	113
	7.7	Discus	sion	113
		7.7.1	Effect of T_2^* relaxation	113

		7.7.2	Setting the noise level and tSNR	114
		7.7.3	Insights from the different scenarios	114
		7.7.4	Limits to the study	115
			Why off-resonance effects are not considered in this study	115
			Restricting to single echo GRE fMRI	116
		7.7.5	Extending the simulator	116
		7.7.6	Exploring the effect of tuning acquisition and reconstruction togethe	er117
		7.7.7	More than a fMRI simulation tool	118
	7.8	Conclu	usion	118
8	Goir	ng deep	per with Plug and Play Priors	121
	Con	text and	d summary	121
	8.1	Introd	uction	122
	8.2	Plug-a	and-Play algorithms for MRI	123
		8.2.1	MR image reconstruction	123
		8.2.2	Plug-and-play algorithms	123
		8.2.3	Complex-valued denoisers for MRI	124
	8.3	Propo	sed approach	124
		8.3.1	Denoised dataset generation	124
		8.3.2	Unsupervised preprocessing of multicoil data	125
		8.3.3	Proposed algorithm	125
		8.3.4	Preconditioning matrices	126
	8.4	Exper	imental results and discussion	127
		8.4.1	Simulated non-Cartesian acquisition	127
		8.4.2	Training of the denoising prior	127
		8.4.3	Image reconstruction results	127
			Preconditioning boosts image quality	127
			Generalization to other acceleration factors	128
			Stability of PnP methods and further acceleration	128
	8.5	Conclu	usion & future work	129
9	Plug	g-and-F	Play reconstruction for 3d non-Cartesian fMRI data	131
	Con	text and	d summary	131
	9.1	Introd	uction	132
	9.2	Mater	ials and Methods	132
		9.2.1	Acquisition simulation setup	132
		9.2.2	Reconstruction with PnP prior	134
		9.2.3	Neural network architecture	135
	9.3	Exper	imental results	136
		9.3.1	Benchmark of MR image reconstruction methods	136
		9.3.2	Training	136
		9.3.3	MR image reconstruction results	136
	o	9.3.4	Statistical analysis of fMRI sequences	137
	9.4	Conclu	usion and Perspectives	138
10	Gen	eral co	nclusion	141
	10.1	Contri	ibutions and limitations	141

Contents

	10.2	Perspectives	142
		10.2.1 Denoising functional MRI data	142
		Additional testing of local-low-rank denoising methods	142
		Explore new denoising methods	143
		10.2.2 fMRI Reconstruction	143
		Perform a full benchmark of fMRI reconstruction methods	143
		Data augmentation for deep learning with SNAKE	144
		Jointly optimizing acquisition and reconstruction for fMRI using	
		SNAKE	144
		Improving and extending SNAKE	144
		10.2.3 Deep-learning for fMRI reconstruction	145
		10.2.4 Towards in-vivo utilization	145
Α	Inve	rse problems and convex optimization essentials	147
	A.1	Inverse Problems and Convex Optimization	147
		A.1.1 Introduction to inverse problem	147
		A.1.2 Convex analysis essentials	149
	A.2	Optimization algorithms	151
		A.2.1 Gradient Descent	152
		A.2.2 Splitting methods	152
		Forward-Backward Splitting	153
		Acceleration	153
		Backward-Backward Splitting	153
		Primal-Dual algorithm	154
		A.2.3 Alternating Direction Method of Multipliers (ADMM)	154
		Acceleration	155
		A.2.4 More than two functions	155
R	Δnn	endix for Chapter 5	157
5	R 1	BrainWeb Phantom	157
	B.1 R 2	Estimation of the regularization parameter using SURE	158
	B.2 B.3	Temporal SNR maps for the three scenarios	158
	D.0	remporar or and the unce seenands	100

Résumé en Français

Motivation et Contexte

L IMAGERIE par résonance magnétique fonctionnelle (IRMf) est une technique largement utilisée pour étudier le fonctionnement du cerveau. Elle est basée sur l'effet BOLD (Blood Oxygenation Level Dependent), qui est un indicateur de l'activité neuronale. Les données de l'IRMf consistent en une série temporelle de volumes 3D, acquis avec une résolution temporelle de l'ordre de la seconde et une résolution spatiale de l'ordre du millimètre. L'origine des signaux IRM fonctionnels provient du contraste BOLD [Oga+90; Kwo+92; Ban+92] et du couplage neurovasculaire qui lie l'activité neuronale à la variation de la proportion de désoxyhémoglobine et d'oxyhémoglobine dans le sang. Heureusement, ces changements sont visibles à l'aide d'un protocole standard d'imagerie par résonance magnétique. En particulier, le contraste pondéré en T_2^* est sensible aux propriétés magnétiques du sang, et le signal BOLD se manifeste par une augmentation transitoire allant jusqu'à 1-5 % de l'intensité du signal IRM dans les régions activées par rapport à la ligne de base.

Comparée aux deux autres modalités de neuro-imagerie non invasives largement utilisées que sont l'électroencéphalographie (EEG) et la magnétoencéphalographie (MEG), l'IRMf offre une excellente localisation spatiale, mais ne permet pas de mesurer directement l'activité neuronale ni d'obtenir une résolution temporelle. Néanmoins, la modalité d'imagerie IRMf a été popularisée depuis son premier développement méthodologique dans les années 1990 et a bénéficié de l'amélioration du matériel et des méthodologies IRM (champ plus élevé, imagerie parallèle, acquisition accélérée), elle est maintenant largement utilisée par les neuroscientifiques pour répondre à des questions fondamentales sur le cerveau sain et pathologique, la conscience et la cognition.

La collecte des données IRMf n'est que la première étape d'un long processus de traitement visant à extraire les informations pertinentes sur le fonctionnement cérébral, soit au repos, soit dans un contexte de stimulation spécifique. Ce processus peut être résumé comme suit:

- 1. Acquisition des données brutes dans le domaine spatial de Fourier (espace k), pendant que le cerveau est stimulé;
- 2. Reconstruction d'une série d'images (ou de volumes);
- 3. Prétraitement de ces séries (individuellement ou par lots) pour compenser les sources connues d'artefacts (comme le mouvement et les effets d'inhomogéneité du champ statique);
- 4. Analyse statistique pour corréler les variations du signal avec le paradigme expérimental.

Chacune de ces étapes a été progressivement améliorée au cours des trois dernières décennies et continuera de l'être.

Si l'IRMf a bénéficié des améliorations de l'acquisition et de la reconstruction de l'IRM anatomique, elle subie également des mêmes limitations et sources de bruit et ajoute à cela une dimension temporelle. Elle souffre d'un compromis quadruple entre la résolution spatiale, la résolution temporelle, le rapport signal sur bruit (RSB) et la couverture du cerveau.

Au cours des dernières décennies, les études d'IRMf ont été menées sur un grand groupe de sujets (comme le Human Brain Project [Ela+21] ou les cohortes ABCD [Cas+18]) et, contraintes par ce compromis, elles visent généralement une résolution spatiale relativement faible (*e.g.* 3mm isotrope), dans le but de minimiser la variabilité inter-sujets et de mener des analyses au niveau du groupe, ce qui conduit à une compréhension du cerveau humain à l'échelle macroscopique. Cependant, pour approfondir notre compréhension de fonctions cérébrales spécifiques, des études à haute résolution sont menées de plus en plus souvent, mais elles se limitent à un nombre restreint de sujets [Pin+20] et tirent parti des avancées des très hauts champs magnétiques [Maz+23; Zha+23], tout en surmontant les défis associés [VP21]. Poussé à ses limites, cela conduit à réduire l'acquisition à une petite partie du cerveau, comme les colonnes corticales d'une zone spécifique [Law+19; YHU08]. La quête d'une couverture complète du cerveau à l'échelle mésoscopique (résolution spatial sous millimétrique, avec une résolution temporelle inférieure à la seconde) se poursuit.

Pour atteindre cet objectif, de nouvelles méthodes d'acquisition et de reconstruction ainsi que des outils d'analyse appropriés doivent être développés.

En outre, l'essor récent de la reconstruction basée sur l'apprentissage profond et les réseaux de neurones dans l'IRM anatomique a montré des résultats prometteurs dans la réduction du temps d'acquisition avec une qualité de reconstruction d'image améliorée [Ram+22; RC23]. L'application de ces méthodes à l'IRMf en est encore à ses balbutiements, en raison de la complexité et de la taille des données, et du manque d'ensembles de données supervisés de référence pour la phase d'entraînement de ses réseaux de neurones. De plus, la reconstruction basée sur l'apprentissage profond pourrait être sujette à un surajustement et à des hallucinations, ce qui invaliderait les résultats de l'analyse statistique. Il est donc crucial de développer des méthodes pour vérifier la qualité des images reconstruites et de comprendre les limites de ces méthodes.

Contributions

Le but de cette thèse est de participer à cet effort, en étudiant de nouvelles méthodes pour les premières étapes de la chaine de traitement de l'IRMf, acquisition et reconstruction, et de proposer des outils pour les explorer et les valider, ouvrant la porte à une optimisation conjointe de ces deux aspects. Les principales contributions de notre travail peuvent être classées comme suit :

Parangonnage des méthodes de débruitage pour l'IRMf

Les données d'IRMf reconstruites acquises à très haut champ (7 Tesla et au-delà) et à haute résolution restent bruitées, en raison du compromis entre le RSB et la résolution spatio-temporelle. Le bruit se présente essentiellement sous la forme d'un bruit blanc

gaussien additif (bruit thermique) [TPW11], et la dimension temporelle qui fournit des informations hautement redondantes peut être exploitée pour débruiter les séries temporelles en 3D.

Toutefois, le niveau de bruit n'est pas nécessairement homogène spatialement. Les travaux déclencheurs de NORDIC [Viz+21] et la famille plus large des méthodes de débruitage de rang spatialement faible (Local-Low-Rank, LLR) déjà utilisée dans l'imagerie pondérée en diffusion (DWI) [Moe+21; Ver+16] ont suscité un grand intérêt dans la communauté de l'IRMf comme moyen d'« abaisser le bruit de fond thermique » des données d'imagerie de l'IRM fonctionelle. Cependant, l'approche empirique décrite a suscité notre intérêt pour l'utilisation de méthodes mathématiquement plus optimales [GD17; GD14] afin de débruiter davantage les données.

Nous avons donc comparé les différentes méthodes de débruitage LLR pour les données IRMf. Nous avons également soulevé la question de l'ordre relatif avec la routine de correction du mouvement généralement utilisée en premier lieu dans le prétraitement de l'analyse de l'IRMf [Com+23b].

La génération de donnée IRMf de référence dans l'espace k, pour le développement de nouvelles méthodes de reconstruction.

Les méthodes de débruitage ne fonctionnent qu'après la reconstruction des images et ne conviennent pas si les données ne peuvent pas être reconstruites. En outre, l'évaluation comparative de ces méthodes a révélé la nécessité d'un outil de comparaison approprié des techniques de reconstruction et de prétraitement de l'IRMf.

À cette fin, nous avons développé "SNAKE" un simulateur d'IRMf de bout en bout, qui modélise le problème de l'IRMf, depuis la spécification des activations cérébrales en tant que réponses BOLD jusqu'à la synthèse de données artificielles de référence dans l'espace k puis la simulation du processus d'acquisition IRMf accélérée passant par un sous-échantillonnage 3D et la prise en compte d'un modèle de décroissance en T2* du signal IRM le long de chaque trajectoire [CVC24a].

SNAKE est un logiciel libre et peut générer efficacement des données à haute résolution spatio-temporelle dans l'espace k et permettre le développement de nouvelles méthodes d'acquisition et de reconstruction d'images optimisées conjointement. Avec SNAKE, nous avons notamment démontré l'efficacité de la stratégie de démarrage à chaud pour la reconstruction de l'IRMf lors de l'utilisation d'un modèle d'acquisition dynamique. En ayant accès aux données de référence, nous pouvons calculer la matrice de confusion et évaluer la qualité de l'image et la performance statistique, en terme de sensibilité et de spécificité de détection, et les relie l'une à l'autre.

Vers une reconstruction basée sur l'apprentissage profond

Enfin, forts des connaissances acquises grâce à SNAKE et à l'évaluation comparative des méthodes de débruitage, nous avons décidé d'intégrer l'étape de débruitage directement dans la méthode de reconstruction, selon une approche dite « Plug-and-Play ». Nous avons d'abord démontré la faisabilité de cette approche dans le contexte de l'IRM anatomique, comme résultats préliminaires avant d'étendre cette approche à l'IRMf.

Aperçu du manuscrit

- **Partie 1** Introduction to functional MRI présente les éléments préliminaires nécessaires à l'introduction du reste du manuscrit.
 - Chapitre 1 Functional MRI: Spying on the brain with magnets introduit les principes fondamentaux de l'IRM fonctionel, le contrast BOLD et les methodes classique d'acquisition en IRMf. Nous y discuterons également comment les expériences IRMf sont conduite en pratique, et quelles sources de bruit les affectent.
 - **Chapitre 2** Getting the fMRI signal faster rappelle comment accélérer l'acquisition des scan IRMf en utilisant l'imagerie parallèle, non-Cartésienne et l'échantilonnage compressé.
- **Partie 2** Reconstructing functional MRI images s'appuyant sur la partie précédente, en se focalisant sur le problème de la reconstruction en IRMf.
 - **Chapitre 3** *Image reconstruction methods for fMRI* donne un aperçu de l'état de l'art pour la reconstruction des données IRMf et introduit également un cadre unifiée pour leur interprétation.
 - Chapitre 4 *Deep Learning for MRI and functional MRI* continue dans cette approche, en considérant l'utilisation des méthodes d'apprentissage profond pour la reconstruction, et décrit également les limites de ces approches pour l'IRMf.
- Partie 3 Methodological developments contient les principales contributions scientifique de ce manuscrit.
 - Chapitre 5 MRI-NUFFT: Doing non-Cartesian MRI has never been easier est le premier chapitre de contribution, et propose un logiciel de calcul de transformée de fourier Non-Uniforme adapté à l'IRM de manière efficiente.
 - **Chapitre 6** Benchmarking local low-rank denoising methods for fMRI compare différentes méthodes de débruitage pour l'IRMf, dans l'idée première de générer des données de réference de haute résolution.
 - **Chapitre 7** No Ground-Truth ? Let's build some présente le simulateur SNAKE, ses hypothèses presents the SNAKE simulator, the underlying modeling hypothesis and implementation details.
 - Chapitre 8 Going deeper with Plug and Play Priors propose d'utiliser les methodes plug-and-play pour la reconstruction d'IRM anatomique.
 - **Chapitre 9** *Plug-and-Play reconstruction for 3d non-Cartesian fMRI data* applique ces méthodes plug-and-play auy cas de l'IRMf et valide leur utilisation grâce au simulateur SNAKE.
 - Chapitre 10 *General conclusion* résume les principales contribution de cette thèse, et décris les perspectives de recherches qu'elles ouvrent.

General Introduction

Context and Motivation

F^{UNCTIONAL} Magnetic Resonance Imaging (fMRI) is a widely used technique to study brain function. It is based on the Blood Oxygenation Level Dependent (BOLD) effect, which is a proxy for neural activity. fMRI data consists of a time series of 3D volumes, acquired with a spatial and temporal resolution in the respective order of millimeters and seconds. The origin of the functional MRI signal arises from the BOLD contrast [Oga+90; Kwo+92; Ban+92] and the neurovascular coupling which links the neuronal activity to the variations of the proportion of deoxyhemoglobin and oxyhemoglobin in the blood. Fortunately, these changes are visible using standard MR scanning protocol. Notably, The T_2^* weighted contrast is sensitive to the magnetic properties of blood, and the BOLD signal witnesses as a transitory increase up to a few percents in the MR signal intensity in the activated brain regions compared to the baseline.

Compared to the two other widely used non-invasive neuro-imaging modalities electroencephalography (EEG) and magnetoencephalograpy (MEG), fMRI offers great spatial localization, but lacks the direct measurement of cerebral activity and temporal resolution. Nonetheless, the fMRI imaging modality has been popularized since its first methodological development in the 1990s and benefited from the improvement of MRI hardware and technologies (higher field strength, parallel imaging, accelerated acquisition), it is now heavily used by neuroscientists to answer fundamental questions about the healthy and pathological brain, and uncover mechanisms behind cognition and consciousness.

Collecting the fMRI data is only the first step of a long processing pipeline to extract relevant information. This pipeline can be summarized as follows

- 1. Acquisition of the raw data in the spatial Fourier domain (k-space), while the brain is stimulated or at rest;
- 2. Reconstruction of a series of image (or volumes);
- 3. Preprocess of fMRI time series (individually or in batch) to compensate for known sources of artifacts (like motion and off-resonance effects);
- 4. Statistical analysis to correlate signal variations with the experimental paradigm.

Each of these steps has been gradually improved across the past three decades, and they will keep on improving.

If fMRI benefited from the improvements of anatomical MRI acquisition and reconstruction, it also suffers from the same limitations and noise sources and adds a temporal dimension. It suffers from a four-way trade-off between spatial resolution, temporal resolution, signal-to-noise ratio (SNR), and brain coverage.

General Introduction

In the past decades fMRI studies have been conducted on a large group of subjects (such as the Human Brain Project [Ela+21], the UK Biobank initiative [Byc+18], or the ABCD cohorts [Cas+18]) and constrained by this trade-off, they typically use low to medium spatial resolution (*e.g.* 3mm isotropic), with the intent to minimize inter-subject variability and conducts group-level analysis, leading to a global scale understanding of the human brain. However, to deepen our understanding of specific brain functions, high-resolution studies are conducted more and more often, but restrict themselves to a limited number of subjects [Pin+20] and leverage the advances of higher field strengths [Maz+23; Zha+23], while overcoming the associated challenges [VP21]. Pushed to its boundaries, this leads to reduce the acquisition to a small portion of the brain, like cortical columns of a specific area [Law+19; YHU08]. The quest for complete coverage of the brain at a mesoscopic scale (*i.e.* submillimetric spatial resolution and sub-second temporal resolution), is still uncharted territory.

To achieve this goal, new acquisition and reconstruction methods as well as suitable analysis tool need to be developed.

Moreover, the recent rise of learning based reconstruction methods in anatomical MRI has shown promising results in reducing the scan time while maintaining quality [Ram+22; RC23]. However, the application of these methods to fMRI is still in its infancy, due to the complexity size and dimensionality (4D) of the data, and the lack of supervised datasets for training. Additionally, current self supervised learning based reconstruction mnethods in fMRI operate slicewise and as such are not able reach high spatio-temporal resolution. As the training dataset are usually not fully sampled, these techniques could be prone to over-fitting and hallucinations, which would invalidate the results of statistical analysis. Therefore, it is crucial to develop methods to verify the quality of the reconstructed images, and to understand the limits of these methods.

Contributions

The goal of this PhD thesis is to participate in this effort, by studying new methods for the early steps of the fMRI pipeline, data acquisition and image reconstruction, and propose tool to explore and validate them, opening the door for joint optimization of these two aspects. The main contributions of our work can be sorted out as follows:

Benchmarking Denoising Methods for fMRI

Current high resolution (at 1mm-isotropic and submillimetric resolution) fMRI images at ultra-high field (7 Tesla and beyond), remain noisy, as a result from the trade-off between SNR and spatio-temporal resolution. The noise is essentially arising as an additive white Gaussian noise, commonly referred to as thermal noise [TPW11], and the temporal dimension that provides highly redundant information can be leveraged to perform denoising of the 3D time-series.

However, the noise level is not necessarily spatially homogeneous. The seminal work of NORDIC [Viz+21] and the broader family of Local-Low-Rank denoising methods already tested in Diffusion Weighted Imaging (DWI) [Moe+21; Ver+16] have shown great interest in the fMRI community as a way to "lower the thermal noise floor" of fMRI images. Yet, the empirical approach outlined in NORDIC sparked our interest in employing more advanced mathematically optimal methods [GD17; GD14] to enhance the denoising of fMRI data. Consequently, we compared various LLR denoising methods for the fMRI data. Additionallu, we examined the relative position of these denoising methods with the motion-correction routine typically used as a first step in fMRI analysis preprocessing[Com+23b].

Getting ground truth fMRI in k-space, for developing new reconstruction methods.

Previous denoising methods only work after the image reconstruction, and are not suitable if the data cannot be reconstructed already. Furthermore, the benchmarking of these methods revealed the need for a proper comparison tool of fMRI reconstruction and preprocessing techniques.

To this end, we developed "SNAKE" an end-to-end fMRI simulator, that models the fMRI process. It encompasses everything from defining brain activity pattern in space and time (i.e. BOLD responses), up to the acquisition of k-space data under an extended Fourier model, taking T_2^* signal decay into account [CVC24a].

SNAKE is open-source and can efficiently generate high spatio-temporal resolution data in k-space, falicitating the development of new jointly optimized acquisition and image reconstruction methods. In particular, with SNAKE we demonstrated the efficiency of a warm-start strategy for fMRI reconstruction when using dynamic acquisition pattern. Access to ground-truth data allows us to compute confusion matrix metrics, to evaluate both image quality and statistical performance, and to determine how they relate to each other.

Towards deep learning based reconstruction

Finally, equipped with the knowledge from SNAKE and the benchmarking of denoising methods, we decided to integrate the denoising step directly in the reconstruction method, as a so called "Plug-n-Play" approach. We first demonstrated the feasibility of this approach in the context of anatomical MRI, as preliminary results to extend this approach to fMRI.

Thesis Outline

- **Part 1** Introduction to functional MRI provides the essential ingredients that are necessary to understand the rest of this work.
 - Chapter 1 Functional MRI: Spying on the brain with magnets introduces the basic principles of fMRI, the BOLD contrast, and the classical acquisition methods for fMRI data. We will also briefly discuss how fMRI experiments are conducted in practice, and how different noises sources affect them.
 - **Chapter 2** *Getting the fMRI signal faster* presents how to accelerate the acquisition of fMRI scans, using parallel imaging, non-Cartesian imaging and compressed sensing.
- Part 2 Reconstructing functional MRI images builds on the previous part to focus on the reconstruction of fMRI data
 - **Chapter 3** *Image reconstruction methods for fMRI* gives an overview of the existing advanced method for fMRI reconstruction and provides an unified framework for their study.
 - **Chapter 4** Deep Learning for MRI and functional MRI goes a step further by considering the use of learning-based reconstruction methods, and the limitation for their applications for fMRI images reconstruction.
- **Part 3** Methodological developments contains the main scientific contributions of this thesis.
 - Chapter 5 *MRI-NUFFT: Doing non-Cartesian MRI has never been easier* is the first contribution chapter, and proposes efficient computation for non-Cartesian imaging.
 - **Chapter 6** Benchmarking local low-rank denoising methods for fMRI compares and evaluates different denoising methods for fMRI, with the first intent to produce high quality reference data.
 - Chapter 7 No Ground-Truth ? Let's build some presents the SNAKE simulator, the underlying modeling hypothesis and implementation details.
 - Chapter 8 Going deeper with Plug and Play Priors presents the preliminary results of the application of the Plug-and-Play approach for MRI data, outlining their potential use in functional MRI.
 - Chapter 9 *Plug-and-Play reconstruction for 3d non-Cartesian fMRI data* goes a step further, and presents the first application of Plug-and-Play methods for fMRI validated using the simulator SNAKE.
 - Chapter 10 *General conclusion* summarizes the main contributions of this thesis, and outlines the perspectives for future work.

List of publications

Articles under revisions for publication in Peer-Reviewed Journals

- P.-A. Comby, A. Vignaud, and P. Ciuciu. SNAKE-fMRI: A Modular fMRI Data Simulator from the Space-Time Domain to k-Space and Back. Version 2. Apr. 12, 2024. URL: http://arxiv.org/abs/2404.08282 (visited on 09/01/2024). In revision to Imaging Neuroscience.
- Pierre-Antoine Comby, Guillaume Daval-Frérot, Caini Pan, Asma Tanabene, Léna Oudjman, Matteo Cencini, Philippe Ciuciu, and Chaithya Gr. *MRI-NUFFT: Doing Non-Cartesian MRI Has Never Been Easier*. Version 2. Nov. 2024. URL: https://hal.science/hal-04775262 (visited on 12/12/2024). in revision to Journal of Open Source Software.

Accepted papers in Peer-reviewed Conferences

- Pierre-Antoine Comby, Matthieu Terris, Alexandre Vignaud, and Philippe Ciuciu. *Plug-and-Play Reconstruction for 3D Non-Cartesian fMRI Data*. Mar. 2025. URL: https://hal.science/hal-04993532. Submitted to EUSIPCO 2025
- P.-A. Comby, B. Lapostolle, M. Terris, and P. Ciuciu. *Robust Plug-and-Play Methods for Highly Accelerated Non-Cartesian MRI Reconstruction*. Version 2. Oct. 2024. URL: https://hal.science/hal-04759015. Accepted to ISBI 2025.
- Z. Amor, P.-A. Comby, C. Le Ster, A. Vignaud, and P. Ciuciu. "Non-Cartesian Non-Fourier FMRI Imaging for High-Resolution Retinotopic Mapping at 7 Tesla". In: 2023 IEEE 9th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP). Dec. 2023, pp. 201–205.
- P.-A. Comby, Z. Amor, A. Vignaud, and P. Ciuciu. "Denoising of FMRI Volumes Using Local Low Rank Methods". In: ISBI 2023 International Symposium on Biomedical Imaging. Apr. 18, 2023.

Abstracts in Peer-reviewed Conferences/ Papers in Workshops

- Z. Amor, P.-A. Comby, P. Ciuciu, and A. Vignaud. "Achieving High Temporal Resolution Using a Sliding-Window Approach for SPARKLING fMRI Data: A Simulation Study". In: *ISMRM Annual Meeting.* Singapore, May 2024.
- P.-A. Comby, A. Vignaud, and P. Ciuciu. "SNAKE-fMRI: A Modular fMRI Simulator from the Space-Time Domain to k-Space Data and Back". In: *ISMRM Annual Meeting, (in Press).* Singapore, 2024. Awarded the 2nd best abstract award from the reproducibility study group.
- P.-A. Comby, G. Daval-Frerot, Chaithya Gr, A. Vignaud, and P. Ciuciu. "MRI-NUFFT: An Open Source Python Package to Make Non-Cartesian MR Imaging Easier". In: *ISMRM Annual Meeting, (in Press)*. Singapore, 2024.
- P.-A. Comby, Z. Amor, A. Vignaud, and P. Ciuciu. "Benchmarking Local Low Rank Denoising Methods for Task-Based fMRI Data Analysis". In: *ISMRM 2023 Annual Meeting.* Toronto, June 2023.

Software Contributions

To achieve the methodological developments of this Ph.D. work, several open-source softwares have been developed:

Logo & Name	Description
MRI-<i>nufft</i> MRI-NUFFT	A unified interface for efficient forward model computations of non-Cartesian MRI. It is used in Chapter 5, 7, and 8.
SNAKE	A simulator for fMRI data in k-space. Described extensively in Chapter 7.
Python Sparse data Analysis Package - IMRI Plugin PySAP-fMRI	Variational methods for reconstructing fMRI time-series. Used in Chapter 7, based on the methods of Chapter 3
patch-denoise	Denoising of fMRI images data using patch-based methods. Used in Chapter 6.
hydra-callbacks & brainweb-dl	A collection of callbacks for the hydra framework and download interface for the BrainWeb Dataset. Used in Chapter 7
benchmark-mri-pnp	A benchmark of Plug-and-Play methods for MRI reconstruction (see Chapter 8).

 Table 1: Software developed during this thesis

Other activities

Student supervision. I supervised *Marie Doggo* a Bachelor student in a double degree in Mathematics and Medicine from Université Paris Cité from April to August 2022, we worked on the impact of denoising on the brain decoding for the fMRI image collected in the IBC project. I also supervised Benjamin Lapostolle, from École Polytechnique, he worked on our first implementation of the Plug-n-Play approach for MRI data, from March to August 2024, paving the way to the results presented in Chapter 7.

Teaching. I was a teaching assistant for the course "Informatique Embarquée" at IUT de Cachan, I gave a practical lesson on how to do C programming on an embedded device, and reminded how timers work, what a pointer is, and to put a semicolon at the end of line.

Reviews. I had the opportunity to review manuscripts for the following journals: *IEEE Transactions on Medical Imaging, IEEE Transactions on Computational Imaging* and *Magnetic Resonance in Medicine*, for which I was also an open source code reviewer.

Communication. I took great pleasure in communicating my results as well as the ones of our team in the form of scientific posters and in making some of my doodles into logos for the software we developed together.

Notation

In this manuscript, the mathematical notation follows general typographical conventions.

\mathbf{Symbol}	Description
\mathbb{N}	The set of Natural numbers.
\mathbb{R}	The set of Real numbers.
α	Lower-case letters denote scalars.
x	Lower-case bold letters denote vectors. In particular $r, k \in \mathbb{R}^3$ represent the spatial position in image domain and k-space, respectively.
A	Upper-case bold letters denote matrices with exceptions for some vectors, notably \boldsymbol{B}_0 and \boldsymbol{M}_0 .
$A^{ op}$ or $x^{ op}$	Transposed matrix (\mathbf{A}) or vector (\mathbf{x}) .
$oldsymbol{A}^H$ or $oldsymbol{x}^H$	Hermitian matrix (\mathbf{A}) or vector (\mathbf{x}) .
$\ oldsymbol{A}\ _{fro}$	Frobenius norm of a matrix defined as $\ \mathbf{A}\ _2 = \sqrt{\sum_{i=1}^m \sum_{j=1}^n A_{i,j} ^2}$.
$\ oldsymbol{x}\ _2$	ℓ_2 -norm of a vector defined as $\ \boldsymbol{x}\ _2 = \sqrt{\boldsymbol{x}^H \boldsymbol{x}} = \sqrt{\sum_{i=1}^m x_i ^2}.$
$\ m{x}\ _W$	Weighted ℓ_2 -norm of a vector defined as $\ \boldsymbol{x}\ _{\boldsymbol{W}} = \ W^{1/2}\boldsymbol{x}\ _2 = \sqrt{\boldsymbol{x}^{\top}\boldsymbol{W}\boldsymbol{x}}.$
$\ m{x}\ _1$	ℓ_1 -norm of a vector defined as $\ \boldsymbol{x}\ _1 = \sum_{i=1}^m x_i $.
$\ oldsymbol{x}\ _0$	ℓ_0 pseudo-norm of a vector defined as $\ \boldsymbol{x}\ _0 = \sum_{i=1}^m 1_{x_i \neq 0}$
$\angle x, x $	The phase and the magnitude of the complex number $x = x e^{i\angle x}$
$\nabla_x f$	Gradient of the function f w.r.t. the variable \boldsymbol{x}
$\operatorname{dom} f$	Domain on which the function f is defined
$\operatorname{prox}_f(x)$	Proximal operator of the function f : $\operatorname{prox}_{f}(x) = \arg\min_{u} f(u) + \frac{1}{2} u - x _{2}^{2}$
$\mathcal{N}(oldsymbol{\mu}, oldsymbol{\Sigma})$	Multivariate normal distribution with mean μ and covariance matrix Σ .

Part I

Introduction to functional MRI

You know, what these people do is really clever. They put little spies into molecules and send radio signals to them, and then they have to radio back what they are seeing — Niels Bohr

> Vous laissez pas embobinez, il va vous rembobiner. — Yvain, Kaamelott

Chapter 1

Functional MRI: Spying on the brain with magnets

Outline

1.1	Source of fMRI Signal: The BOLD contrast	16
	1.1.1 The neurovascular coupling	16
	1.1.2 The Hemodynamic Response Function model	17
1.2	Getting the fMRI signal	17
	1.2.1 Nuclear Magnetic Resonance: the Bloch model	17
	1.2.2 From Bloch equations to the Fourier acquisition model	21
	1.2.3 Pulse Sequences for fMRI	24
1.3	fMRI in practice	27
	1.3.1 the fMRI experiment	27
	1.3.2 Noise and artifacts sources and their mitigation strategies	29
	1.3.3 fMRI Analysis with the General Linear Model	31
1.4	Conclusion	33

THE human brain accounts for 2% percent of body mass, but consumes 20% of our resources in glucose and oxygen. Its 86 billion cells, which are a few micrometers in diameter, can span several dozens of centimeters, are connected between each other and share information through electrical and chemical signals. At our macroscopic scale, they are the source of our consciousness, thoughts, and action. Trying to understand brain functions is a quest that is at the intersection of many disciplines: biology, psychology, medicine, and philosophy.

Yet, Physicist have developed an invaluable tool at the intersection of two of their favorite fields: quantum mechanics and electromagnetism, namely the Magnetic Resonance Imaging (MRI) Scanner.

In this chapter, we will first briefly describe the effects that link brain activity to changes in physical properties that an MRI signal acquisition can sample, through the neurovascular coupling and the BOLD effect. Then, we will present the basic principles of nuclear magnetic resonance, the Bloch model, and the acquisition of the MRI signal using a standard pulse sequence: the gradient recall Echo Echo planar Imaging GRE EPI. Lastly, we will focus on how fMRI experiments are conducted in practice, and how different noise sources and artifacts affect the quality of the acquired data. We will also present the general processing pipeline of fMRI data that is usually carried on to overcome these limitations. For more detailed explanations, the reader is invited to consult the following books on MRI [Nis96; Ped21; Bro+14] and fMRI [UUB15; UK23; Ash19] Furthermore, multiple Ph.D. theses written in the past in the MIND and METRIC team provide a summary of MRI Physics [Rad23; Dav22b; Amo24]. They have been invaluable in my understanding the phenomenon of magnetic resonance imaging and its application to neuroscience.

1.1 Source of fMRI Signal: The BOLD contrast

1.1.1 The neurovascular coupling

The brain is the central hub for processing sensory inputs and sends commands to the rest of the body. In order to receive the necessary nutrients to fulfill its mission, it is tightly connected to the vasculature system, and most of the exchange between blood vessels and the neurons occurs in the gray matter. Therefore, it has been observed (as early as the 19th century) that cerebral activity is linked to local changes in blood flow, volume and oxygenation in the cortex. The relation between these blood¹ dynamics and the underlying neural activity is called the neurovascular coupling. The intrinsic details of this coupling are still an active area of research (See notably the balloon model [Fri+00; BF97] and how it translates in change of magnetic properties [UMU09; UK23]). In a nutshell, an increase in cerebral activity will have three cumulative effects:

- An increase in oxygen consumption (CMRO₂): this results in a decrease in oxyhemoglobin concentration (O_2Hb) and an increase in the concentration of deoxyhemoglobin (Hb).
- An increase in cerebral blood flow (CBF).
- An increase in cerebral blood volume (CBV).

CBF and CBV overcompensate for oxygen consumption and, in general, lead to an increase in the ratio $\frac{[O_2Hb]}{[Hb]}$



Figure 1.1: The neurovascular coupling and the hemodynamic response function.

Deoxyhemoglobin (*Hb*) is paramagnetic (as opposed to oxyhemoglobin (O_2Hb which is diamagnetic), and distort the local magnetic field. Thus, the increase in the ratio $\frac{[O_2Hb]}{[Hb]}$

¹"Haemon", $Ai\mu\omega\nu$, blood in ancient Greek

provokes a change in the acquired MR signal. This change is called the **Blood Oxygen** Level Dependent (BOLD) contrast, and is the basis of the functional MRI signal.

Remark. The BOLD effect is essentially witnessed as a transient increase of a few percents in the gray matter (GM). There are also recent works on BOLD signal detection in white matter (WM) [Wan+23; Zu+24], but this research direction is still in its infancy. Furthermore, neuronal activity can be tracked using other modalities such MEG and EEG (with great temporal resolution) but their lack of spatial resolution hinders their use to determine the exact localization of these effects.

1.1.2 The Hemodynamic Response Function model

The connection between the BOLD signal and neural activity does not occur instantaneously, as the dynamics of neurovascular coupling, observed over a period of several seconds, lags behind cerebral activity [Glo99]. Under initial assumptions, this sequence of events is represented by a linear time-invariant system, with the impulse response of this system referred to as the Hemodynamic Response Function (HRF). The canonical HRF (represented in Figure 1.1) is usually modeled as a difference of two Gamma functions, but other models have been proposed, extending its parametrization to include spatial information and more [Ciu+03; Vin+14; BVC13].

1.2 Getting the fMRI signal

We saw that the neurovascular coupling provides a proxy mapping from cerebral activity to local modification of magnetic properties. In the following we will leverage this effect to measure these changes using an MRI system.

1.2.1 Nuclear Magnetic Resonance: the Bloch model

Magnetization vector

Getting information from the brain (or any other organ) using an MRI system relies on the following three aspects:

- 1. Source: Hydrogen protons ${}^{1}H$ constitute 80% of our bodies (mainly in water, but are also present in lipids).
- 2. Polarization: When placed in a strong uniform magnetic field B_0 , a sample of protons will behave like a tiny rotating magnet² with strength M aligned with the B_0 field, and precess around the direction of B_0 at the so-called Larmor frequency $\omega_0 = \gamma B_0$, with γ the gyromagnetic ratio³. In steady state, we have $||M|| = M_0 \propto B_0$. This group of protons will be called the magnetization cell hereafter.

 $^{^2{\}rm This}$ classical mechanical analogy is used here for brevity. In fact, this phenomenon arises from the combination of quantum properties.

 $^{^{3}\}gamma = 42.58 \,\mathrm{MHz} \,\mathrm{T}^{-1}$ for the hydrogen proton.

3. Excitation and Relaxation: As M_0 is small and aligned with B_0 , it is not feasible to measure it directly; Instead, it must first be excited by an external field B_1^+ , oscillating at the Larmor frequency and orthogonal to B_0 . This will tip M onto the transverse plan, and once B_1^+ is turned off, M will return to its steady state. This change can be witnessed using the RF coils that captures the signal from the transverse plane.

The complex interactions between protons and their environment can be summarized by the Bloch equation.⁴

$$\frac{\partial \boldsymbol{M}}{\partial t} = \gamma \boldsymbol{M} \times \boldsymbol{B} + \frac{\left(M_0 - M_z\right)\hat{\boldsymbol{z}}}{T_1(\boldsymbol{r})} - \frac{M_x \hat{\boldsymbol{x}} + M_y \hat{\boldsymbol{y}}}{T_2(\boldsymbol{r})},\tag{1.1}$$

with $\mathbf{M}(\mathbf{r},t)$ is the magnetization at position \mathbf{r} , $\mathbf{B} = [\mathbf{B}_{1,x}^+(\mathbf{r},t), \mathbf{B}_{1,y}^+(\mathbf{r},t), \mathbf{B}_0]$ the magnetic field apply to the sample. $\mathbf{B}_{1,xy}^+$ is the exciting field, orthogonal to \mathbf{B}_0 , and is typically applied for a short period of time δT .

 $T_1(\mathbf{r})$ and $T_2(\mathbf{r})$ describes the longitudinal relaxation time (it takes T_1 seconds for M_z to reach back 68% of the original M_0 steady state magnetization) and the transverse magnetization (it takes T_2 seconds for M_x and M_y to fall of 37% from their original value). T_1 and T_2 are spatially varying, and are characterizing the local environment.

Remark. For the Bloch model to be valid, it is required to place the size of our *magnetization cell* at a sweet spot where we have enough protons to use the magnetization formalization and ignore quantum effects, but the region described by position \boldsymbol{r} is small enough for \boldsymbol{B}, T_1, T_2 to be considered constant at position \boldsymbol{r} . This grouping of protons is called an *isochromat* in the NMR literature.

In practice, the transversal relaxation is faster due to inhomogeneities in the B_0 field inside the isochromat and the observed time T_2^* is rather defined as:

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \gamma \delta \mathbf{B}_0 \,. \tag{1.2}$$

In particular, different biological tissues have different values of T_1 and T_2 . Different types of tissues (*e.g.* in the human brain: white matter, gray matter, cerebrospinal fluid, blood, dura matter, etc.) will therefore experience different relaxation times.

To retrieve this information, MR physicists manipulate the magnetic field $\boldsymbol{B}(\boldsymbol{r},t)$ in space and time using \boldsymbol{B}_0 and \boldsymbol{B}_1^+ and measure the response of magnetization using antenna (or "coils"). From this signal, they are able to get a contrasted image that will depend on T_1 and T_2 . The full magnetic field $\boldsymbol{B}(\boldsymbol{r},t)$ can be expressed as follows:

$$\boldsymbol{B} = \begin{bmatrix} B_{1,x}^+(\boldsymbol{r},t) & B_{1,y}^+(\boldsymbol{r},t), & B_0 + \boldsymbol{G}(t) \cdot \boldsymbol{r} + \Delta B_0 \end{bmatrix}.$$
(1.3)

 $G(t) \cdot \mathbf{r}$ is a linearly varying magnetic field on top of B_0 . With it, the spatial location of the magnetization is encoded in its precessing frequency $\omega_0(\mathbf{r},t) = \gamma(B_0 + \mathbf{G}(t) \cdot \mathbf{r})$. $\Delta B_0(\mathbf{r},t)$ represents imperfection in the main B_0 field, such imperfections are always present in practice due to a change in magnetic susceptibility⁵.

⁴Effects such as diffusion or chemical shifts [Bro+14] are not considered in this work.

⁵See the PhD dissertation of Guillaume Daval-Frérot [Dav22a] for more insight around this issue.

Remark. In practice, in a MRI machine, the magnetic field B is created using a set dedicated coils to create the following components:

- B_0 in the order of Tesla (typical values are 1.5, 3, and up to 7 and 11.7T).
- $G(t) \cdot r$ is the Gradient Field, obtained using three coils inserted in the bore each generating a spatially varying component (G_x, G_y, G_z) in the order of 100 mT/m.
- B_1^+ is created using an antenna/coil placed close to the organ to be scanned and produces a field in the order of a few μ T.

Excitation

Solving the Bloch model, while B_1^+ is effective, is beyond the scope of this thesis, but for our purposes, the main outcome of its application is the flipping of the magnetization to an angle α :

$$\alpha = \gamma \int_0^{\delta T} B_1^+(t) \, \mathrm{d}t \,. \tag{1.4}$$

Relaxation

By setting $M_{xy} = M_x + iM_y$ and if the $B_{1,xy}^+$ field has been turned off we can determine the evolution of the longitudinal and transverse magnetization. A graphical representation is given in Figure 1.2a:

$$\begin{cases} M_z(\mathbf{r},t) = M_z(\mathbf{r},0)e^{-t/T_1(\mathbf{r})} + M_0(1-e^{-t/T_1(\mathbf{r})}) \\ M_{xy}(\mathbf{r},t) = M_{xy}(\mathbf{r},0)e^{i\omega_0 t}e^{-i\Delta\phi(\mathbf{r},t)}e^{-t/T_2^*(\mathbf{r})}, \end{cases}$$
(1.5)

where $\Delta \phi$ represents the phase accrual due to the evolution of $\boldsymbol{G}(t) \cdot \boldsymbol{r}$:

$$\Delta\phi(x,y,z,t) = \gamma \left(\int_0^t G_x(\tau) x \mathrm{d}\tau + \int_0^t G_y(\tau) y \mathrm{d}\tau + \int_0^t G_z(\tau) z \mathrm{d}\tau \right) = \gamma \mathbf{k}(t) \cdot \mathbf{r} \,. \tag{1.6}$$

Through $\mathbf{k}(t) = \int_0^t \mathbf{G}(\tau) d\tau$, we encode how the spatial location of the sample evolves as a function of the Larmor frequency.

Steady state magnetization

If the exciting field B_1^+ is repeated with a period TR (foreshadowing § 1.2.3), we can use the continuity hypothesis to determine steady-state magnetization.

If TR is long enough, such that the transverse magnetization has been null out due to T_2^* decay⁶, but remains short enough so that the longitudinal magnetization is not fully recovered⁷, we can derive a steady-state magnetization by stating that in (1.5) $M_z(t=0) = M_z(t=TR) \cos(\alpha)$ (the magnetization at t=TR just being flipped by a new pulse with an angle α). We note this value M_z^0 :

 $^{^{6}}$ This can be enforced using so-called spoiling gradients for some time before the TR.

 $^{^{7}}T_{1}$ is typically 10 to 50 times larger than T_{2}^{*} in brain tissues.

Chapter 1 Functional MRI: Spying on the brain with magnets

$$M_z^0 = M_z(0) = M_z(TR)\cos(\alpha)$$
(1.7)

$$M_z^0 = M_0(1 - e^{-TR/T_1}) + M_z^0 \cos(\alpha) e^{-TR/T_1}$$
(1.8)

$$M_z^0 = \frac{M_0(1 - e^{-TR/T_1})}{1 - e^{-TR/T_1}\cos(\alpha)}$$
(1.9)

Once B_1^+ is applied and assuming that no residual transverse magnetization was left from the previous excitation, we have the following, given (1.7):

$$M_{xy}(0) = M_z^0 \sin(\alpha) = \frac{M_0(1 - e^{-TR/T_1})}{1 - e^{-TR/T_1} \cos(\alpha)} \sin(\alpha).$$
(1.10)

Remark. Maximizing $M_{xy}(0)$ can be done by choosing $\alpha = \arccos(e^{-TR/T_1})$, called the *Ernst angle.* In the context of functional magnetic resonance imaging, we would therefore set up α according to the T_1 value of the gray matter.



(b) Free Induction decay for two different T_2^* values. The oscillation at ω_0 is only shown for the shortest T_2^* .

Figure 1.2: Relaxation of the magnetization vector and the Free Induction Decay (FID) signal.

Neurovascular coupling and NMR

As we mentioned earlier, neurovascular coupling links cerebral activity to a local change in blood oxygenation, CBF and CBV (see also [KO12] for a deeper analysis of this relation).

Through the magnetic properties of deoxyhemoglobin, this changes overall translate into a decrease in T_2^* at the location containing blood vessels submitted to the neurovascular coupling [UK23; Jin+06]. More precisely, if we model this change as $\frac{1}{T_{2,BOLD}^*} = \frac{1}{T_2^*} + \frac{1}{\delta T_2^*}$ (with $\delta T_2^* < 0$) then, the contrast between two magnetization vectors with and without the BOLD effect can be expressed as:

$$\frac{\Delta M_{xy}}{M_{xy}} = \frac{M_{xy,B} - M_{xy}}{M_{xy}} = \frac{e^{-t/T_{2,B}^*} - e^{-t/T_2^*}}{e^{-t/T_2^*}} = e^{-t/\delta T_2^*} - 1$$
(1.11)

$$\frac{\Delta M_{xy}}{M_{xy}} \simeq -t/\delta T_2^* \tag{1.12}$$

Thus, as T_2^* is modified by cerebral activity, we will observe an increase in signal value due to the BOLD effect.

Remark. This result may lead to the idea that one should sample the magnetization as late as possible to get the best contrast. In practice, this is not feasible because we are limited by the Signal-to-Noise Ratio (SNR), which decrease along with the transverse relaxation and the need for temporal resolution. Furthermore, the HRF model introduces a temporal variation to δT_2^* , and the above development considered that δT_2^* is constant for TR. Due to the typical value of TR (20 ms to 100 ms) compared to the duration of the HRF (in the order of 20 s to 30 s) this can be considered a valid hypothesis. It should also be noted that T_2^* values measured in tissues decrease as B_0 increases as modeled in Eq. (1.2), [vdZwa+09]. More complex BOLD models exist (notably accounting for its effect on the phase of the signal) [UMU09; HU20; CA12], but they are not considered in this work for the following reasons: First, the BOLD effect is in it majority still visible in this change of magnitude. Then, the usual fMRI analysis is conducted on magnitude-only data, discarding the phase information. Finally, they would require to add more terms to the Bloch equation (1.1), notably taking into account blood flow dynamic and variable proton density.

1.2.2 From Bloch equations to the Fourier acquisition model

Capturing the signal

With the help of electromagnetism laws [Bro+14] we can obtain the evolution of all the magnetization by placing a coil near the sample and measuring the voltage at its ends. After numerous mathematical manipulations, such as applying demodulation to get rid of the $e^{i\omega_0 t}$ term resulting from the base precession due to B_0 , we are left with a measured signal defined as:

$$s(t) = \int_{FOV} M_{xy}(\boldsymbol{r}, 0) B_{1,xy}^{-}(\boldsymbol{r}) e^{-t/T_{2}^{*}(\boldsymbol{r})} e^{-i\gamma \boldsymbol{k}(t) \cdot \boldsymbol{r}} \mathrm{d}\boldsymbol{r}$$
(1.13)

and using (1.10) we get:

$$s(t) = \int_{FOV} M_0(\boldsymbol{r}) \frac{(1 - e^{-TR/T_1(\boldsymbol{r})})\sin(\alpha)}{1 - e^{-TR/T_1(\boldsymbol{r})}\cos(\alpha)} B_{1,xy}^-(\boldsymbol{r}) e^{-t/T_2^*(\boldsymbol{r})} e^{-i\gamma \boldsymbol{k}(t)\cdot\boldsymbol{r}} \mathrm{d}\boldsymbol{r}$$
(1.14)

$$\simeq \int_{FOV} x(\boldsymbol{r}) \mathcal{S}(\boldsymbol{r}) e^{-i\gamma \boldsymbol{k}(t) \cdot \boldsymbol{r}} \mathrm{d}\boldsymbol{r} \,. \tag{1.15}$$

Eq. (1.14) is one of the most important equations in this thesis. Several observations can be drawn from it:

- $B_{1,xy}^-$ is the transverse component receiving field (complex-valued, using the same formalism as for M_{xy}). It is specific to each coil and depends on its geometry, position and the load applied by placing an object nearby, it really represents the *point of view* of the coil. As such, it is also called the *sensitivity coil profile* S. It can be shown that the sensitivity profile of the coil is dependent on the coil geometry and the sample itself and is slowly varying in space.[Web16]
- If no gradient is applied, $\mathbf{k}(t) = 0$ and the signal decay is known as *Free Induction Decay* (FID), and is represented in Figure 1.2b.
- As it describes the *forward model* of the MR system, we will find it again in the simulation setup of Chapter 7.
- Eq. (1.15) neglects the T_2^* decay, and with no temporal dependency in the integral, it is the Fourier transform of $\boldsymbol{x}(\boldsymbol{r})\mathcal{S}(\boldsymbol{r})$ (the magnetization viewed from the coil point of view), sampled at the locations $\boldsymbol{k}(t)$. This will be the starting point for Chapter 3.
- The MR physicist can control the magnetization (through RF pulse B_1^+) and spatially encode this information (using G(t) and (1.6)) to obtain a signal that depends on the local relaxation times T_1 and T_2^* and the spatial distribution of the magnetization. These timely managed manipulations are known as *pulse sequences* (they will be covered in more detail in § 1.2.3).

Image Formation and k-space encoding

The signal s(t) is sampled using an Analog-to-Digital Converter (ADC) at a sampling rate δt called the *dwell time* (in the order of 1 µs to 10 µs). Thus, we acquire a set of values $y_{\ell,i} = s_{\ell}(t_i)$ for the ℓ -th coil, at the frequencies $\mathbf{k}_1, \ldots, \mathbf{k}_N$ from the instants t_1, \ldots, t_N .

In this discrete setting, we can rewrite the acquisition process as

$$y_{\ell,i} = \sum_{j=1}^{N} x(\boldsymbol{r}_j) \mathcal{S}_{\ell}(\boldsymbol{r}_j) e^{-i\gamma \boldsymbol{k}_i \cdot \boldsymbol{r}_j} + \epsilon_{\ell,i}$$
(1.16)

$$\boldsymbol{y}_{\ell} = \mathcal{F}_{\Omega} \mathcal{S}_{\ell} \boldsymbol{x} + \boldsymbol{\epsilon}_{\ell} \tag{1.17}$$

where \mathcal{F}_{Ω} is the Fourier transform, restricted to the frequencies Ω , and ϵ_{ℓ} is an additive noise (typically considered Gaussian)

The number of frequencies to be sampled is determined by the field of view (FOV) and the target image resolution. According to the Shannon-Nyquist theorem, to faith-fully retrieve an image of size $(N_x, N_y, N_z \text{ of a FOV} (FOV_x, FOV_y, FOV_z))$, the spacing between two frequencies must be $\Delta k_u = \frac{1}{FOV_u}, \forall u \in \{x, y, z\}$. Moreover, the maximum frequency range $[-K_u^{\max}; K_u^{\max}]$ is $K_u^{\max} = \frac{N_u}{2FOV_u}$.

The set of all acquired spatial frequencies is called the k-space. However, the decay T_2^* , as well as physical constraints on the gradients and the RF pulse, limit the portion of sampled k-space during a single read-out window. Thus, the excitation-relaxation cycle (also called a *shot*) is repeated, and different k-space locations are collected during each shot. The way and order in which the k-space is traversed is called a *k-space trajectory*.

Two examples of 2D k-space trajectory are presented in Figure 1.3. A Cartesian single shot, called Echo Planar Imaging (EPI), implements a zigzag sampling scheme on the grid from the top left to the bottom-right corner, while an in-out spiral starts from the left hemifield, converges to the center of the k space through clockwise rotation, touches the center of k-space at its mid-journey, and then makes a U-turn to implement the opposite movement in the counter-clockwise sense and reaches the right. The lower values of k represent the center of the k-space, which contains the low-frequency image information such as contrast (and most of the signal energy), while the higher values of k represent the periphery of the k-space, which contains the high-frequency information (edges and details).

Remark. The time at which the center of k-space is acquired is called the *echo time TE*, i.e. when the average phase of the magnetization is zero, and the signal amplitude is thus maximal. Along with the TR, and the flip angle α (§ 1.2.1), they are key parameters defining the image contrast.



Figure 1.3: Examples of two classical k-space sampling pattern used in fMRI and their associated gradients profiles. Top: Echo Planar Imaging with k-space trajectory (left) and piecewise constant gradient waveforms (right). One gradient is played at a time. Bottom: In-out Archimedean spiral trajectory (left). Sinusoidal gradient waveforms in quadrature (right). These plots were produced using MRI-NUFFT. The frequencies are normalized to [-0.5, 0.5].

1.2.3 Pulse Sequences for fMRI

Echo mechanism and Bold sensitivity

At the end of § 1.2.1, we saw that the BOLD effect is witnessed by the change in T_2^* , and mainly in the gray matter. A good pulse sequence for fMRI should thus aim to maximize the sensibility to this effect. The contrast between an image is given by the difference of their signal at echo time TE [UK23], and we can express the BOLD contrast as the difference between baseline and activated signal, as presented in Figure 1.4. From this we can determine that the optimal TE to maximize the BOLD sensitivity is to use a value close to the T_2^* of the (activated) gray matter. Moreover the BOLD effect is more visible at higher field strength [vdZwa+09], but with a shorter T_2^* . This is for instance shown in Figure 1.4: The peak BOLD sensitivity is 50% greater at 7T than at 3T, when sampled at the right TE (30 ms at 7T, 52 ms at 3T)



Figure 1.4: BOLD sensitivity and its impact on TE, adapted from [UK23]. The BOLD sensitivity is maximal at the T_2^* value of the activated gray matter. The BOLD sensitivity has been magnified 10 times for visualization purposes

Remark. At higher field strength, the length of the acquisition also has to be shorten, as the decay of the BOLD sensitivity is faster than at 3T[UK23]. This means that the overall acquisition window for good signal quality is shorter at higher field strength, and that one has to find efficient ways to sample the k-space in a shorter time. However, the SNR is also higher at higher field strength, which helps the overall detection of the BOLD signal.

Finally, it is also worth noting that the BOLD effect can also be witnessed in T_2 contrast, obtained using a Spin-Echo (SE) sequence [PMN10; Bux+04]. T_2 contrast is more sensitive to the capillaries effect than to the big vessels contained in a voxel. However, such sequences are rarely used in practice, as they require higher energy deposition and longer acquisition times [CV23].
Constraints on the pulse sequence

In the context of fMRI, there are several constraints on the pulse sequence:

- Sensitivity to the BOLD effect. As shown earlier, the effect of the BOLD signal to the base signal is proportional to the echo time TE. As presented in Figure 1.4 the optimal TE has to be tuned to the T_2^* value of the gray matter, where the BOLD effect is expected, to get the maximum sensitivity.
- **Isotropic resolution.** To obtain a true representation of the image, isotropic resolution is required to avoid sensitivity to specific alignment in the brain, and manipulate the information at the same resolution in every direction. Furthermore, a high spatial resolution (e.g. 1mm isotropic) is desired to increase the specificity of the BOLD signal detection.
- **Temporal Resolution.** To be sufficiently sensitive to the evoked or spontaneous brain activity and spatial fluctuations of the HRF, a high temporal resolution is also desirable for the development of fMRI pulse sequences.
- **Brain coverage.** Outside very specific studies, whole brain coverage is also required to visualize complete brain activity and perform for instance functional connectivity studies that unveil how different distant brain region communicate together. obtain the complete cerebral response to a given stimulus or to probe brain activity at rest and decipher the organization of intrinsic resting-state functional networks.
- **Field strength.** The field strength of the MR system is also a key factor in the detection of the BOLD signal. Higher field strengths have several advantages:
 - (i) They increase the base SNR of the acquisition as $M_0 \propto B_0$. And experimentally $\text{SNR}_{voxel} \propto B_0^{\alpha}$ (with $\alpha \geq 1.5$, and experimentally estimated as $\alpha = 1.65$ [PSS16]).
 - (ii) The BOLD effect is stronger [vdZwa+09], (but overall the decay of T_2^* is also faster requiring shorter TE), and the spatial specificity of the BOLD signal increases.

However, it also comes with two major drawbacks: At higher field strengths, inhomogeneities and imperfections are more detrimental and need specific corrections at both the acquisition and image reconstruction stages. First, B_1^+ suffers from a shorter wavelength due to the higher Larmor frequency. This is typically overcome with parallel transmission during RF pulse delivery [Gra+17b]. Second, a large-scale field B_0 inhomogeneity is present at the air-tissue interfaces, thus inducing geometric distortions and signal loss in the image [Dav22a]. This must be compensated for during image reconstruction.

Remark. The constraints presented above could also be formulated as trying to maximize the BOLD *Contrast-to-Noise-Ratio* (CNR_{BOLD}) while maintaining high spatial and temporal resolution for enhanced specificity and sensitivity. Increasing the CNR means on one hand getting a lower noise level, which is typically obtained by moving to higher magnetic field, increasing the voxel size (more signal per unit of volume), and improved coil design, or on the other hand getting stronger effect (increased δT_2^*) which is also achieved by reaching higher magnetic field (Figure 1.4). In other words, using the formalism of [Duy12] and the values of [PSS16], the CNR can be defined as:

$$CNR = \frac{BOLD \text{ sensitivity}}{\text{noise level}} \propto SNR \cdot \frac{\delta T_2^*}{T_2^*} \propto \Delta \boldsymbol{r} \cdot B_0^{1+1.65}$$
(1.18)

where $\Delta \mathbf{r}$ represent the volume of a voxel. This model (and the above quadratic gain of CNR with an increased B_0 justifies the will to go to build more powerful MRI systems, and overcome all the challenges it brings. SNR and its impact on the data will be further discussed in § 1.3.2, and along the rest of this work.

The most common pulse sequence used in fMRI is the *Gradient Recalled Echo Echo Planar Imaging* (GRE-EPI) sequence. It is a fast sequence that is sensitive to the BOLD effect and can be used to acquire a full brain volume in a few seconds, by scanning a full (or partial) plan of k-space after a single excitation.

Gradient Recalled Echo

Gradient Recalled Echo (GRE) sequence is the most common type of pulse sequence used for fMRI. It is a fast sequence that is sensitive to the BOLD effect and can be used to acquire a full brain volume in 2 seconds or less. The main idea for the GRE sequence is first to excite the magnetization with a pulse that tips the magnetization at an angle α and then applies a gradient to move to a starting point in the k-space. Then the ADC is turned on and the signal is acquired while the k-space is traversed.

Remark. The name of the GRE sequence arises from how the signal is acquired: The echo of the signal is obtained at *TE*, in the center of k-space, *i.e.* when the gradient are null.

The effectiveness of the GRE sequence for fMRI arise from the choice of the long readout window (like EPI, presented below).

Echo Planar Imaging

In Echo Planar Imaging (EPI), a single readout of a k-space plane is executed during a single shot, and the k-space is traversed using a sequence of gradients. First, a sliceselective gradient is applied with the RF pulse to excite a slice of the brain. Then, a phase encoding gradient is applied to move to a specific location in the k-space, and the signal is acquired on a line-by-line basis using the last gradient dimension as a readout direction. The slice selection mechanism means that the acquisition is done on a 2D Fourier plane, and repeated to obtain a full 3D volume. Due to its fast sampling and robustness to different artifacts, 2D-EPI is the most common sequence used in fMRI, and various improvements have been made to it, such as multi-band EPI (also called Simultaneous Multi Slice or SMS) [Bar+15] and 3D EPI [Pos+10]. See [LSte+19] for a comparison between the two approaches and their limitations and Figure 1.5 for an illustration of the 3D EPI pulse sequence diagram.



Figure 1.5: Pulse sequence diagram for a GRE 3D EPI. The whole volume is excited by a single RF pulse (pink), then gradients (green, blue, reds) play the trajectory to traverse the k-space along the prescribed trajectory, while the ADC is activated to acquire the received signal. The center of the k-space plane is reached at the time TE. This is repeated at a period TR to acquire all k-space planes of the volume.

Remark. Multi-Echo fMRI is also feasible: After a single excitation pulse, the k-space is sampled multiple times (typically 2 or 3 times) after a single pulse excitation and each time the center of k-space is sampled we get a new echo (with exponentially decaying intensity). The data from each echo can be combined to partially overcome B_0 inhomogeneities, but this come at the expense of a lower temporal resolution for a given TR and extra processing [Kun+17; Gül+23].

1.3 fMRI in practice

1.3.1 the fMRI experiment

fMRI experiments can be classified in two main categories:

- *task-based* fMRI: A participant is submitted to controlled sensory stimuli or cognitive tasks. The participant's cerebral activity recorded via the BOLD signal will be correlated with the stimuli events.
- *resting-state* fMRI: the subject remains at rest and is asked to do "nothing". Here, the goal is to extract information on spontaneous brain activity, notably to extract intrinsic functional brain networks, also called resting-state networks, either in a time-locked manner or using dynamic methods.

Chapter 1 Functional MRI: Spying on the brain with magnets

Remark. In the context of this thesis, we will focus on task-based fMRI. However, the proposed approaches could be extended to resting-state fMRI with minimal effort.

In any case, the intensity of the BOLD effect is quite minimal (1-5% peak increase compared to the baseline).

fMRI equipment

An MRI machine is a complex piece of engineering, and we saw how the pulse sequence drives the different components of the machine. Figure 1.6 shows, how a volunteer will undergo the experiment. For task-based fMRI, visual stimuli are presented to the subject using a projector, the volunteer can provide feedback through push-buttons. All this equipment must be "MR"-compatible to not disturb the acquired signal. The main magnetic field B_0 and the gradients coils are behind the tunnel wall. The RF coil for transmission and reception are embedded inside a helmet placed on the head of the subject, which is stabilized with foam pads. The volunteer is installed on a bed that can be moved inside the bore.



Figure 1.6: Myself, getting set up for a mock fMRI experiment at the occasion of the 11.7T Iseult Scanner press release. The experiment actually happens inside the bore. Credit: *Alain Jocard/AFP*.

fMRI software pipeline

Toolboxes and software for fMRI data analysis are numerous, and the choice of the software will depend on the specific needs (i.e., statistical analysis) of the study.

Since most of the fMRI data have been acquired with a Cartesian sampling pattern, the reconstruction is not much more complicated than a plain inverse FFT, and the subsequent processing takes place in the image domain. Moreover, such reconstructed data is directly an output of the scanner in the form of DICOM image files.

There are several software packages widely used in the fMRI community, such as FSL [FSL17], SPM [Fri07], AFNI [Cox96], and their development dates back to the infancy of fMRI in the 1990s. With the increase in fMRI cohorts and the need for reproducibility, new standards have emerged, such as the BIDS convention [Gor+16] for data organization and its preprocessing is now routinely performed using fMRIprep [Est+19], which allows for parallel processing of the data and has become the standard *de facto* for fMRI preprocessing.

The main goal of all these preprocessing steps is to clean the data from noise and artifacts (described in the next session), to boost the extraction of the BOLD signal and then to perform some statistical analysis (*e.g.* fixed or random effect general linear model at the subject and/or group level) or to project the fMRI images onto an existing anatomical template for group-level analysis.

1.3.2 Noise and artifacts sources and their mitigation strategies

In practice, an fMRI run does not contain only a clean BOLD signal, but also several sources of noise such as thermal and physiological ones. To quantify the quality of the signal, MR Physics relies on the signal-to-noise-ratio (SNR) spatially and temporally.

Remark. Among different communities, there exists different definition of the signal-tonoise ratio. For the MR-physicist, since he cannot access noiseless reference data, the signal is usually an average over a dimension or a neighborhood (*e.g.* where signal is the strongest) and the noise level is the standard deviation over the same dimension, or a different neighborhood (*e.g.* the background of the image). This differs from the definition of electrical engineers, who use instead a ratio of energy levels (*i.e.* average of square signal, and max variance).

Thermal Noise and SNR

Outside the ideal model, the acquisition of $y_{\ell}(\mathbf{k}_i)$ is corrupted by thermal noise arising from the underlying Brownian motion of the spin of ¹H as well as the electronic preprocessing chain. In general, the SNR in the k-space SNR_k can be shown to be proportional to the bandwidth of the acquisition signal (*i.e.* the dwell time of the ADC).

In the image domain, and considering that the k-space is fully sampled, the SNR can be shown to be in 3D:

$$\text{SNR}_{image} \propto B_0^{\alpha} \Delta_x \Delta_y \Delta_z \sqrt{N_x N_y N_z \delta t}$$
 (1.19)

In other words, the SNR scales with its volume, longer acquisition time, and higher field strength. Thus, fMRI acquisition setup are bound by a 4-way trade-off between spatial resolution, temporal resolution, brain coverage, and desired SNR level.

Physiological Noise

The temporal dimension of the fMRI signal means that other temporal variations than BOLD are also captured. In particular, breathing, heart beat, and rigid head motion are also visible in the signal. With improved SNR_{img} , these variations are also better

Chapter 1 Functional MRI: Spying on the brain with magnets

captured and limit the so-called tSNR, the SNR along the time dimension (as opposed to the image dimensions of SNR_{img}), as shown in Figure 1.7. The proposed relation is:

(1.20)



Figure 1.7: SNR-tSNR mapping. Original figure from [Tri+05]. The Spatial SNR plotted along the x-axis and called SNR_0 was modulated using different voxel sizes and field strengths.

All these subject-induced variations can be mitigated by using external devices and procedures such as cardiac and respiratory cycle monitoring, head motion tracking, and retrospective motion correction. Nevertheless, there will still be some residual physiological noise left, than can be accounted for in the fMRI analysis.

Field Inhomogeneities

The last major effect that degrades the BOLD signal is large field inhomogeneities. If microscopic inhomogeneities are considered in the signal model by relaxation of T_2^* (and are the source of BOLD contrast), macroscopic inhomogeneities in the main magnetic field B_0 can lead to geometric distortions and loss of signal in the image.

The source of these inhomogeneities is twofold: First, the field produced by the main magnet is never perfectly homogeneous (despite the best efforts of the vendors); Second, the field is perturbed by the presence of the subject itself which induces magnetic susceptibility changes, notably at the air-tissue interface in the sinuses, ear canal and the skull. In addition, these inhomogeneities can vary over time (instabilities in the gradient system, eddy currents, and subject-related changes such as breathing or heart beat).

The static effects can be corrected using a field map, acquired using a specific pulse sequence, and the dynamic effects can be corrected using a navigator echo, acquired during the acquisition itself (which reduces the temporal resolution of the acquisition), or using a field-monitoring system.

If this additional information is available, retrospective reconstruction can be used to compensate for field inhomogeneities [Amo+23; SNF03; Amo+24a].

Remark. For EPI, field inhomogeneities manifest themselves as geometric distortions. Thus, their correction typically consists in two consecutive acquisitions with opposite phase encoding gradients, and post processing using the TOPUP algorithm in FSL, to estimate and correct the distortion between the two images.

At ultra-high field (7 Telsa and more), field inhomogeneities also concern the RF pulse, and its B_1^+ field. Indeed, the higher Larmor frequency results in a shorter wavelength, smaller than the typical human head. In this case, the energy deposition of the RF pulse is not homogeneous, resulting in discrepancy in the effective flip angle, and ultimately a loss of contrast. The most elegant solution to overcome B_1^+ is certainly parallel transmission (pTX) [Gra+17b], which uses multiple transmission coils (potentially different from the receiving ones) to send a collection of pulses that will provide a full and homogeneous coverage of the brain.

1.3.3 fMRI Analysis with the General Linear Model

The goal of fMRI statistical analysis is to link cerebral activity (and localization) with a given experimental paradigm. This is usually done by performing a regression analysis between the expected brain response and the measured BOLD signal. The expected brain response is typically obtained by convolving the sequence of onsets of a given experimental condition with the HRF. To account for potential delay in the time-to-peak, additional regressors can be introduced in the design matrix, for instance, the convolution with the first-order temporal derivative of the HRF. For details, see SPM12 Manual or Nilearn Documentation[Nil].



Figure 1.8: Time series picked up in a given voxel that elicits a BOLD signal in response to a stimulus. Here a block design approach is presented, with 3 alternating activating and resting periods, each period lasting $\simeq 20$ s

Multiple effects can impact the measured time series. Without further knowledge on their relation, we can consider a linear additive relationship. Thus, to analyze the potential impact of each factor (also called *regressors*) we can formalize the following general linear model, voxel-wise:

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon},\tag{1.21}$$

where \boldsymbol{Y} is the time series of a given voxel, \boldsymbol{X} is the so-called design matrix, each of its columns represents a regression, *i.e.* a potential contribution to the fMRI signal and thus



Figure 1.9: Example of activation maps obtained for an auditory stimulus. The activating voxels are located in the primary and secondary auditory cortex near the ear canals. From Nilearn documentation

is built from the expected BOLD response. It can also contain additional regressors that account for residual contributions (motion, breathing, and cardiac activity, for example). In Figure 1.8 we provide an example of a block based paradigm, its expected BOLD response (obtained by convolving the event onsets with the HRF), and the measured signal in one voxel where the BOLD response is happening.

For each voxel, we can fit β under a least-square criterion (assuming that ε is an additive white Gaussian noise, if not a additional pre-whitening step can be performed). β is called the effect size parameter. Then, the importance of each effect can be estimated using a statistical test, like the Student's *t*-test⁸, and thresholded at a specific *p*-value for evaluating the significance. Typical value are p < 0.05, p < 0.01, p < 0.001 respectively predicting that there is 5%, 1% or 0.1% chance that the witnessed effect is due to random effects. At the end of our subject-level analysis, we are able to get an *activation map* as the one shown in Figure 1.9.

However, since the test is performed on every voxel in the image, the chance of random occurrence is multiplied by the number of tests. This multiple comparison effect must be taken into account in statistical detection using, for instance, the Bonferroni correction (lowering the value of p, which may then be too conservative) or by controlling for the false discovery rate using the Benjamini-Hochberg method [Ash19].

Remark. Statistical analysis of fMRI data has had its controversy in the past [BMW09] and in the absence of ground-truth, comparing fMRI methodologies remains a complex endeavor.

Moreover, statistical analysis are usually conducted on multiples subjects to get better significance and robustness, merging the data from different subjects requires further processing and analysis, which we will not cover in this work.

⁸with $H_0: \mathbf{c}^\top \boldsymbol{\beta} = 0$, where \mathbf{c} is a contrast vector selecting one or several component of $\boldsymbol{\beta}$



Figure 1.10: fMRI processing chain pyramid. The left side shows the logos and names of tools used in the literature at this step. On the right side are examples of the fMRI data at the various stages of processing. Inside the pyramid we recall the core equations driving each step.

1.4 Conclusion

This chapter introduced the fundamental principles driving the fMRI signal source and its acquisition using an MRI machine. Getting information from the brain using fMRI is a complex process that starts with the acquisition (typically conducted using a GRE-EPI sequence), reconstruction of the data using a Fourier model, some preprocessing steps to boost the signal quality and then the extraction of the BOLD response to a given stimulus using a GLM. As depicted in Figure 1.10, proper acquisition and reconstruction are required foundations on which further processing and analysis is built.

Underlying all these developments is the quest for a better SNR to increase the sensitivity and sensibility of the BOLD response. Yet, the four-way trade-off between spatial resolution, temporal resolution, brain coverage, and SNR is a constant challenge that the MR physicist must face.

Yet, acquisition of fMRI data is only the first step of a long processing pipeline to extract the BOLD signal and its statistical significance according to the experimental setup. Improvements for the acquisition strategies, which is the first step to get better source signal will be presented in Chapter 2.

In Chapter 3 and 4 we will discuss further the reconstruction of fMRI data and its challenges. Furthermore, Chapter 6 will review a novel early-step of the fMRI processing pipeline, to remove thermal noise and improve the SNR of the data. Finally, it should be noted that the model presented in this chapter will also be the starting point for the simulation of fMRI data used in Chapter 7.

Work it harder, make it better Do it faster, makes us stronger More than ever, hour after hour Work is never over — Daft Punk

Chapter 2

Getting the fMRI signal faster

Outline

2.1	Parallel imaging	35
	2.1.1 GRAPPA	36
	2.1.2 SENSE	37
	2.1.3 2D vs 3D imaging	37
2.2	Non Cartesian imaging	38
	2.2.1 Spiral trajectories for fMRI	38
	2.2.2 Radial-Like trajectories for fMRI	39
2.3	Compressed Sensing	40
	2.3.1 Sparsity	41
	2.3.2 Variable Density Sampling	42
	2.3.3 Optimizing the sampling trajectories	43
	2.3.4 Compressed sensing in practice	43
	2.3.5 Combining CS and PI	44
2.4	Accelerating fMRI data acquisition using non-Cartesian compressed	
	sensing	44
2.5	Conclusion	46

T N chapter 1, we described the underlying principle for fMRI acquisition and outlined the necessity to find a compromise between spatial and temporal resolution while maintaining a good signal-to-noise ratio. However, the basic EPI scheme is clearly suboptimal to reach the end-goal of high-resolution fMRI with standard MR hardware [Fei+23] (≤ 1 mm isotropic spatial resolution and 1s temporal resolution) that would be desirable for getting new insights on brain functions. This quest for high-resolution fMRI has driven the development of new hardware and methodologies to sample more information faster, while maintaining good signal quality. Aside from getting higher field strength to boost the SNR and overcoming the associated challenges, efforts can also be made on the acquisition and image reconstruction side, that this chapter will present: Parallel imaging (§ 2.1), non-Cartesian acquisition (§ 2.2) and compressed sensing (§ 2.3).

2.1 Parallel imaging

In Eq. (1.16) we outlined the dependence on the coil sensitivity profile. By placing multiple coils around the organ of interest (*e.g.* the brain), we can obtain multiple

measurements of the same signal and thus additional information about the underlying magnetization distribution of this organ.

This redundancy of information can be leveraged to under-sample the k space, using an *Acceleration Factor* (AF) (in the context of Cartesian imaging, an acceleration factor R means acquiring only every R line). The undersampled data acquired through the set of coils can then be combined to fill the missing information either in k-space (GRAPPA, § 2.1.1) or in the image multicoil domain (SENSE, § 2.1.2).



Figure 2.1: Multi-coil arrays: Each coil captures the MR signal with a unique sensitivity profile determined by its location around the subject. Adapted from Zaineb Amor's dissertation [Amo24; Des+12]

2.1.1 **GRAPPA**

GRAPPA (Generalized Autocalibrating Partially Parallel Acquisition)[Gri+02] works in the k-space domain. In GRAPPA acceleration, the under-sampled data of each coil is processed in the k-space domain, and the missing lines for a given coil are estimated using the information from the neighboring lines from all coils together using a convolution kernel as an interpolator. This kernel is estimated using a set of fully sampled calibration lines, which are localized in the k-space center.

Remark. GRAPPA-like methods have been refined over the years, and their ease of use in the Cartesian setting has been dominant. More recent development such as RAKI [Akç+19] improves on these base principles.

SNR in Parallel Imaging

Reducing acquisition scan time through under-sampling is not coming without drawbacks. The undersampled data remains noisy, and spreading the information on the sampled points to reconstruct missing information will amplify the noise. It can be shown that if one accelerates the acquisition by a factor R, the SNR of the reconstructed image will be impacted as follows:

$$SNR_R = \frac{SNR_0}{g\sqrt{R}} \tag{2.1}$$

where SNR_0 is the original SNR level without accelerating and g is the so-called gfactor, dependent on the coil geometry and the under-sampling pattern. The g-factor measures the noise amplification due to the undersampling process. In particular, this results in a spatially heterogeneous noise level in the image. Mitigation of the g factor map is possible prospectively by choosing a well-designed subsampling pattern such as CAIPIRINHA[Bre+05] or retrospectively using a smoothing regularization term in the reconstruction algorithm to reduce the impact of the noise.

2.1.2 SENSE

SENSE (Sensitivity Encoding) [Pru+99] works in the image domain. The undersampled data of each coil are transformed in the image domain and yield an aliased image because we do not fulfill the Shannon-Nyquist criteria. However, if one has access to the sensitivity map of each coil, then the underlying organ can be visualized accurately by solving a linear system of equations for each voxel, or minimizing the corresponding least square criterion. As SENSE is formulated in the image domain, it also allows for non-Cartesian sampling trajectories and can be combined with Compressed Sensing, as detailed later on.

Remark. SENSE requires the knowledge of the coil sensitivity profiles. Such profiles can be obtained separately by performing a fully sampled very low-resolution acquisition or estimated from the center of the k-space data itself.

2.1.3 2D vs 3D imaging

The everyday fMRI experiment is conducted with a 2D EPI sequence, where the RF pulse excites only a single slice of the volume of interest. However, 2D imaging limits the acceleration opportunities (only along the phase encoding direction, e.g. k_y). Using a 3D excitation pulse, it is also possible to accelerate along the slice partition dimension (*e.g.* k_z) and combine this acceleration with the others in the first two dimensions.

Remark. It is also possible to accelerate 2D acquisition with simultaneous multi-slice acquisition (SMS), where several 2D slices are excited and acquired simultaneously, and disentangle at the reconstruction stage [Bar+15]. This has some limitations (especially in term of energy deposition) as outlined in [LSte+19].

In addition, 3D imaging also allows for more freedom in the design of k-space trajectories, as the latter can be optimized for a specific criterion, such as matching a specific target sampling density while complying with the hardware constraints involving the gradient magnetic fields.

To reach higher acceleration factor, it is however necessary to break free from the Cartesian equispaced undersampling.

2.2 Non Cartesian imaging

By sampling off the Cartesian grid, it is possible to further accelerate the acquisition process in MRI. A wide variety of non-Cartesian sampling patterns have been proposed over the last two decades. A small collection of them is represented in Figure 2.2.

Remark. Note however, that the traditional radial and spiral sampling pattern requires 2π more sampling points (and thus shots) to reach the Nyquist criterion. The acceleration from under-sampling will arise from the combination of non-Cartesian sampling pattern with Parallel Imaging and (in the next section) Compressed Sensing.



Figure 2.2: 3D Non-Cartesian undersampling patterns for MRI: *Top*: Stack-ofspirals [PHA17],Phyllotaxis radial¹ [Pic+11], Yarn ball [SB21]; *Bottom*: TUR-BINE [GMC22], REPI [RMS21] and SPARKLING [Cha+22]. Additional patterns are available in the MRI-NUFFT documentation.²

2.2.1 Spiral trajectories for fMRI

In non-Cartesian settings, the most popular sampling pattern for fMRI is the 2D spiral sampling pattern and its different variations [Kas+22; Eng+18; Glo12].

Compared to the generally used EPI sequence, spiral imaging distinguishes itself by the simultaneous use of the phase and frequency spatial encoding gradients, which allows for a more efficient sampling of the k-space (see Figure 1.3 from previous chapter). The spiral trajectory can be designed to start at the center of the k-space and move outward (spiral-out) or in the opposite direction (spiral-in), an in-out trajectory can also be used [Kas+22;

¹Phyllotaxis radial is not really suitable for fMRI application as it consists of short readout radial spokes.

²Thank you to Guillaume Daval-Frérot for all the hard work on implementing these trajectories and more.

Glo12]. Spiral methods have shown be less sensitive to motion artifacts and provides typically shorter readout times. However, spiral imaging also possesses intrinsic drawbacks, in particular, the off-resonance artifacts, which manifests as geometric distortion in EPI, results in blurring and signal loss for spirals. The non-Cartesian acquisition also requires a non-Cartesian reconstruction, which is more computationally intensive (Which will be addressed in Chapter 5) than the Cartesian one, and has been one of the major roadblock to its adoption in the MRI community. Nonetheless, spiral trajectories have shown great potential for fMRI, with increased tSNR map, as shown in Figure 2.3.

Finally it should be noted that non-Cartesian acquisition provides extra freedom in the sampling pattern, and this can be used to optimize the acquisition for a specific criterion or create more efficient undersampling patterns, as shown in the next section. Similarly the reconstruction can be made more complex as it will be presented in Chapter 3. It also worth mentioning here that one can for instance reconstruct the whole sequence of fMRI volume at once using global priors (such as low-rank+sparse [Chi+15; PHA17])





Figure 2.3: TSNR map comparing EPI and spiral based trajectory, from [Glo12]. The spiral trajectory shows a better tSNR in the frontal and temporal lobes.

2.2.2 Radial-Like trajectories for fMRI

Another popular choice for doing non-Cartesian MRI acquisition is the radial sampling pattern, which has proven to be naturally robust to motion artifact, as it consists of multiple spokes each going through the center of k-space, and the Fourier transform

Chapter 2 Getting the fMRI signal faster

is equivariant to rotation¹. However, radial imaging in its simplest form is inherently slow, as each spokes requires its own excitation pulse for a relative short readout window. Instead, crossbreeding of EPI and radial trajectory is used such as in TURBINE [Gra+17a; GMC22] and REPI [RMS21].



Figure 2.4: TURBINE vs EPI: Reproduced from [GMC22]. Inset highlight resolution difference in the cerebellum (green) and artifacts in the temporal lobes (white).

In particular TURBINE has shown to be a good alternative to EPI as presented in Figure 2.4. However, both TURBINE and REPI have only been used in a limited number of subjects, and more prospective validation is required to reach broader adoption. Moreover, in both case the resolution is limited to $1 \times 1 \times 2$ mm³ (REPI) and $0.8 \times 0.8 \times 2$ mm³ (TURBINE), with a low temporal resolution (2.4s for TURBINE) for whole brain acquisition.

Remark. These non-Cartesian acquisition requires a more complex and challenging reconstruction (the authors of REPI mention a 24hour offline reconstruction for a single fMRI acquisition). We will discuss these aspects in the next section and provides further solution both in term of algorithm (in Chapter 3) and efficient implementation in Chapter 5.

2.3 Compressed Sensing

The development of Compressed Sensing (CS) theory [CRT06] has found a fertile ground in MRI imaging for its application [OUL20; Lus+08]. However, it remains underused in the fMRI community. In a nutshell, CS is a theory that states that if a signal is sparse in a certain domain (e.g. a wavelet basis), then it can be reconstructed from a few measurements collected in the acquisition domain at a rate below the Shannon-Nyquist criterion if these measurements are undersampled in an incoherent manner, which means by minimizing the correlation between the sampling and image domain. This incoherency

 ${}^{1}\boldsymbol{R}\mathcal{F}(x) = \mathcal{F}(\boldsymbol{R}x)$

can be achieved for instance using a variable density sampling (§ 2.3.2) An example of the effect of CS undersampling and reconstruction is given in Figure 2.5.

In the MRI context, the sparsity is usually modeled in the wavelet domain [Mal09] while under-sampling is performed in the k-space domain. In what follows, we formalize the core principles of Compressed Sensing and present how it is implemented in practice.



Figure 2.5: Example Compressed sensing reconstruction in the 2D Cartesian case. Here no parallel imaging is used for simplicity, at AF=4, the compressed sensing reconstruction can retrieve the ground truth faithfully, in contrast to the zero-filled reconstruction where no *a priori* is made and a simple adjoint Fourier transform is used. Adapted from PySAP-MRI [Far+20] examples.

2.3.1 Sparsity

Natural images have an inherent structure (*i.e.* they are not purely noise), and even partial knowledge can be used to represent these images in a compact way. More formally, there should exist a set of basis functions $\Psi \in \mathbb{C}^{N \times K}$, such that an image \boldsymbol{x} has a decomposition \boldsymbol{z} : $\boldsymbol{x} = \Psi^H \boldsymbol{z}$, \boldsymbol{z} being a sparse vector, *i.e.* a vector with a few nonzero coefficients. More formally, let us define the following Hilbert spaces:

- $\mathcal{X} = \mathbb{C}^N$ is the image space.
- $\mathcal{Y} = \mathbb{C}^N$ defines the observation space (*i.e.* the k-space in MRI).
- $\mathcal{Z} = \mathbb{C}^K$ is the sparse representation domain.

We now introduce the following operators to move from one (potentially incomplete) representation of \boldsymbol{x} to another:



where \mathcal{F}_{Ω} is the Fourier Transform restricted to the k-space sample locations $\Omega = (\mathbf{k}_1, \ldots, \mathbf{k}_M)$, with $M \leq N$. Then, for $\mathbf{y} = \mathcal{F}_{\Omega} \mathbf{x}$, the CS reconstruction can be formulated as follows:

$$\hat{\boldsymbol{z}} = \operatorname*{arg\,min}_{\boldsymbol{z} \in \mathbb{C}^{K}} \|\boldsymbol{z}\|_{0}$$
 subject to $\boldsymbol{y} = \Omega \mathcal{F} \Psi^{H} \boldsymbol{z}$ (2.3)

$$\hat{\boldsymbol{x}} = \boldsymbol{\Psi}^{H} \hat{\boldsymbol{z}} \tag{2.4}$$

In practice, the ℓ_0 pseudo-norm problem is NP-hard. Thus, it is often replaced by a relaxed version, involving the convex hull of the pseudo-norm, the ℓ_1 norm. The new minimization problem, called the *Basis Pursuit*, is formulated as follows:

$$\hat{\boldsymbol{z}} = \arg\min_{\boldsymbol{z}\in\mathbb{C}^K} \|\boldsymbol{z}\|_1$$
 subject to $\boldsymbol{y} = \mathcal{F}_{\Omega}\Psi^H \boldsymbol{z}$. (2.5)

We will now answer the following two questions:

- 1. How to optimally choose the sampling locations Ω ?
- 2. How is Eq. (2.5) solved in practice?

2.3.2 Variable Density Sampling

A key aspect of CS theory is that we can establish guarantees to perfectly reconstruct \boldsymbol{x} with high probability. This notably depends on how the samples in the measurement space \mathcal{Y} are drawn, in regards to the sparsity transform. To make this more explicit, let $\mathcal{A} = \mathcal{F}_{\Omega} \Psi^{H} = [\boldsymbol{a}_{1}, \ldots, \boldsymbol{a}_{N}] \in \mathbb{C}^{N \times K}$ denote the full measurement operator from the sparse domain to the measurement domain. Then, we can use the following theorem [Rau10; Cha+14], adapted to our notations:

Theorem 2.1. Given the above definitions, let:

- (i) z a sparse vector with s non-zeros entries, such that their sign/phase follows a uniform distribution.
- (ii) $p = p_1, \ldots, p_N$, the probability of sampling the locations $\omega_1, \ldots, \omega_N$.

Define:

$$K(\mathcal{A}, p) = \max_{k \in \{1...n\}} \frac{\|a_k\|_{\infty}^2}{p_k}$$
(2.6)

Sample m independent points following the probability distribution p, such that

$$m \ge CK(\mathcal{A}, p)s \log\left(\frac{N}{1-\eta}\right)$$
 (2.7)

where C is a constant. Then, with probability $(1 - \eta)$, z is the unique solution of (2.5).

This theoretical foundation must be instantiated in the MRI context. The goal is to lower the bound in (2.7), and to minimize the number of measurements m. To do so, while providing guarantees on the reconstruction fidelity $(\eta \rightarrow 0)$ it is thus required to:

- Find a good sparsifying transform Ψ to get a very sparse (small s) representation,
- Optimize the probability of sampling a point to minimize $K(\mathcal{A}, p)$.

In particular, it can be shown that if Ψ is a wavelet transform, the points located in the center of the k-space should be sampled more frequently (they correspond to the non-sparse approximation coefficients in the wavelet decomposition). Thus, optimal CS undersampling in MRI implies *variable density sampling* in the k-space. However, matching the exact distribution is not so straightforward in practice due to the hardware constraints on the maximal gradient amplitude and slew rate. Additionally, the limited duration of each shot affects the segmentation of the k-space acquisition and, eventually, the scan time.

Remark. Suitable probability densities have been studied in the literature [CRC21; Kno+11], following the seminal work of [Lus+08]. However, it is worth mentioning that the feasibility of the trajectory, or at least its optimality, is not necessarily taken into account in these studies, reducing their practical impact. More recently, the optimization of the sampling pattern has gained renewed interest in the context of Deep Learning [Wei+21b; Wan+22; RC23].

An example of a 1D variable density sampling pattern (along the phase encoding direction) in the context of Cartesian MRI is given in Figure 2.5.

2.3.3 Optimizing the sampling trajectories

The extra degrees of freedom provided by sampling off the Cartesian grid allow us to match an optimized target sampling density more closely than simple Cartesian lines while reducing the acquisition time.

However, these possibilities are not limitless. First, hardware and safety concerns impose strong constraints on the gradient waveforms. Second, in gradient echo sequences, the echo time should always be placed at the same time point in the center of k-space to maintain the contrast consistency across shots, The trajectories can thus be optimized for a specific target sampling density while satisfying constraints. This idea is at the core of the development of the SPARKLING class of trajectories [Rad+22; Cha+22; Laz+19].

In addition to hardware constraints, the acquisition of MRI data along non-Cartesian trajectories is also sensitive to field inhomogeneities, and mitigation strategies must be implemented for improved results[Amo+22b; Amo+23; Rad+22]. We now turn our attention to the practical implementation of CS reconstruction.

2.3.4 Compressed sensing in practice

Solvable formulation

There is no closed-form solution to the optimization problem (2.5). Usually, the idea is to transform it into a penalized version, where the weight of the ℓ_1 -norm is controlled by a parameter μ :

$$\hat{\boldsymbol{z}} = \arg\min_{\boldsymbol{z}\in\mathbb{C}^{K}} \frac{1}{2} \| \mathcal{F}_{\Omega} \boldsymbol{\Psi}^{H} \boldsymbol{z} - \boldsymbol{y} \|_{2}^{2} + \mu \| \boldsymbol{z} \|_{1}.$$
(2.8)

This formulation can then be solved using the iterative Forward-Backward algorithm or its variants (see \S A.2.2).

Analysis vs Synthesis

The formalism presented above is often termed the *synthesis* regularization, as optimization lies in the sparse domain. In a nutshell, the image is assumed to be generated from the sparse representation. In contrast, complementary *analysis* regularization performs the optimization directly in the image space and enforces sparsity in the transform domain by applying the transformation within the regularization term. The two approaches are equivalent if the sparsifying transform is a tight frame (*i.e.* $\Psi^{H}\Psi = \nu Id$), which is a tenable assumption in the case of orthogonal wavelet transforms.

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}\in\mathbb{C}^N} \frac{1}{2} \|\mathcal{F}_{\Omega}\boldsymbol{x} - \boldsymbol{y}\|_2^2 + \mu \|\boldsymbol{\Psi}\boldsymbol{x}\|_1.$$
(2.9)

However, the analysis approach has been shown to be more effective for MRI reconstruction when using redundant decompositions (*e.g.* undecimated wavelet transform) are used as sparse transforms [Che+18]

Remark. Wavelet decomposition is a common sparsifying transform for CS-MRI reconstruction, but it is not specifically tuned to the problem. Instead, the *dictionary learning* approach proposes learning the atoms of decomposition in a joint optimization with the reconstruction of the image [Iqb+20; Hua+14; Rav+18; RCF17].

2.3.5 Combining CS and PI

Multi-coil compressed sensing reconstruction for a single MR image can be formulated as follows:

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \frac{1}{2} \sum_{\ell=1}^{L} \| \mathcal{F}_{\Omega} \mathcal{S}_{\ell} \boldsymbol{x} - \boldsymbol{y}_{t,\ell} \|_{2}^{2} + \mu \| \boldsymbol{\Psi} \boldsymbol{x} \|_{1}$$
(2.10)

where \mathcal{F}_{Ω} is the Fourier transform sampled as the locations Ω , \mathcal{S}_{ℓ} is the sensitivity map, \boldsymbol{y}_{ℓ} the k-space data associated with the ℓ -th coil, and $\boldsymbol{\Psi}$ is the wavelet transform. If the sensitivity maps are not known in advance, they can be estimated before the reconstruction (for instance, using ESPIRIT [Uec+14] or by masking the center of k space), or the problem can be formulated as a *calibrationless* reconstruction, in which coil-specific images are reconstructed together and combined afterward. In this case, group sparsity can be enforced along the channel dimension [El +19].

Remark. An alternative to the variational reconstruction presented above is to use DNNbased reconstruction (typically by unrolling an iterative reconstruction algorithm), these approaches will be discussed in more detail in Chapter 4

2.4 Accelerating fMRI data acquisition using non-Cartesian compressed sensing

The temporal dimension in functional MRI can also be leveraged to further accelerate the acquisition (with an according reconstruction procedure). In the context of Cartesian acquisition we already mentioned CAIPI methods, however, akin to the traditional acceleration of SENSE and GRAPPA they rely on coherent undersampling, and remains limited to the Cartesian grid.

In functional magnetic resonance imaging (and more generally in dynamic MRI cases), optimization of the sampling trajectories could also take advantage of the temporal dimension. This remains an active area of research. Coherent temporal aliasing has been proposed in approaches known as CAIPIRINHA [Bre+05], which stick to the Cartesian imaging setting for ease of fMRI image reconstruction. However, these methods have limited acceleration power and do not use the efficiency provided by the compressed sensing framework.

By varying the sampling trajectory across volumes — and with adequate reconstruction strategies, it is possible to further accelerate the acquisition in dynamic MRI.

This is notably the case for cardiac imaging: The CINE method [AE91] leverages the periodicity of the heartbeat cycle to perform the reconstruction of a single average period at high temporal resolution from the incoherent acquisition of multiple low temporal resolution cycles.

Remark. Unlike cardiac magnetic resonance imaging, fMRI reconstructs the full-time series, as the stimuli and the ensuing hemodynamic responses are not necessarily periodic. Furthermore, variations in T_2^* -weighted contrast are more subtle in fMRI. Due to the time-frequency duality, a BOLD signal that will be well localized in space and time on fMRI scans will be hard to delineate in the spatial and temporal frequency domains. The challenging question of *where does the BOLD lie in the k-space?* has not yet received a definitive answer.

Currently, the application of CS in fMRI is still relatively restricted. Initial research by Zaineb Amor [Amo+23; Amo24] at NeuroSpin has demonstrated that integrating non-Cartesian trajectories, ultra-high magnetic field and comprehensive brain acquisition in fMRI is demanding. To optimize computational efficiency and facilitate scan parallelization, each volume was independently reconstructed. This process had to address both static and dynamic field inhomogeneities to reach optimal results. More details are available in § 5.3.

Remark. Another point of concern of using CS for fMRI reconstruction is the impact of the non-linear reconstruction on the statistical analysis. If further theoretical analysis is required on this point, the empirical case has proven to hold true, notably on the gaussian noise hypothesis for the GLM analysis (see the supplementary material of [Amo+24b] for more details).

In the fMRI community, other contributions have paid attention to the acceleration of data acquisition along the temporal dimension with time-varying acquisitions and typical low-rank + sparse regularization terms [Chi+15; GMC22; PHA17]. However, these methods require keeping the full k-space data in memory at all time points, and the reconstruction algorithms used are challenging to tune. A more detailed review of the state of the art for the reconstruction of the whole 4D sequence will be provided in Chapter 3.



Figure 2.6: Multi-coil, non-Cartesian, compressed sensing based reconstruction for a single anatomical MRI volume.

2.5 Conclusion

In this chapter, we presented three different strategies to accelerate MRI acquisition: parallel imaging (§ 2.1), non-Cartesian acquisition (§ 2.2), and compressed sensing (§ 2.3). In particular, the combination of these three strategies has shown encouraging results, comparable with the best setup in Cartesian settings. However, such methods give more leeway in their reconstruction, that is more complex than a plain FFT as presented in Figure 2.6 for a simple frame-by-frame reconstruction, and requires efficient non-uniform Fourier Transform, that will be addressed in Chapter 5. This freedom in how the data could be reconstructed is both a blessing and curse that the MR Physicist delegates to signal and image processing scientists. The realm of variational reconstruction methods will be explored in more detail in Chapter 3. The trendy use of deep learning for the reconstruction will be discussed later in Chapter 4.

Part II

Reconstructing functional MRI images

C'est le temps que tu as perdu pour ta rose qui rend ta rose si importante. — Le petit prince, Antoine de St-Exupery

Chapter 3

Image reconstruction methods for fMRI

Outline

3.1	Compressed Sensing and fMRI	49
	3.1.1 General problem formulation	49
	3.1.2 Frame-wise reconstruction	50
	3.1.3 Time-aware reconstruction	51
3.2	Further extensions	53
	3.2.1 Other explicit priors	53
3.3	Conclusion	54

T^N Chapter 1 and Chapter 2, we detailed the core principle driving the acquisition of (f)MRI data, and introduced the standard methods for reconstructing an MR image from undersampled k-space data. However, functional MRI provides a fourth, temporal, dimension to the data, that can be leveraged for more advanced reconstruction methods.

In this chapter, we will explore the different possibilities for reconstructing fMRI images using the classic compressed sensing based approaches and their extensions. This chapter relies heavily on the mathematical tools introduced in Chapter A, the reader unfamiliar with solving inverse problems using convex optimization is encouraged to refer to this appendix.

3.1 Compressed Sensing and fMRI

3.1.1 General problem formulation

In the context of fMRI, the goal is to reconstruct a set of volumes from k-space measurements collected along trajectories or shots.

Let $\boldsymbol{x} \in \mathbb{C}^{N_t \times K}$, be spatiotemporal data $(K = N_x N_y N_z)$, acquired with an MRI scanner. In practice, we collect measurements in the spatial Fourier domain (k-space) $\boldsymbol{y}_{t,\ell} \in \mathbb{C}^M$ for each time frame $t = 1, \ldots N_t$ and each coil ℓ . Here, we assume that the k-space data are collected over non-Cartesian shots in a variable density manner to meet the CS principles. The fMRI images to be reconstructed are not sparse in the image domain but instead in a transform domain such as wavelets. In addition, in fMRI, we can enforce supplementary prior knowledge regarding the time dimension: The temporal evolution of the image sequence is of low rank. Based on this a priori knowledge, we want to solve the following optimization problem:

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x} \in \mathbb{C}^{N_t \times K}}{\operatorname{arg\,min}} \underbrace{\frac{1}{2} \sum_{t=1}^{N_t} \sum_{\ell=1}^{L} \|\mathcal{F}_{\Omega_t} \mathcal{S}_{\ell} \boldsymbol{x}_{t,\cdot} - \boldsymbol{y}_{t,\ell}\|_2^2}_{f(\boldsymbol{x})}}_{g(\boldsymbol{x})} + \underbrace{\lambda_t \sum_{k=1}^{K} \|\mathcal{T} \boldsymbol{x}_{\cdot,k}\|_1}_{g_t(\boldsymbol{x})} + \underbrace{\lambda_s \sum_{t=1}^{N_t} \|P_s \Psi \boldsymbol{x}_{t,\cdot}\|_1}_{g_s(\boldsymbol{x})} + \underbrace{\lambda_a \|P_a \Psi \boldsymbol{x}\|_*}_{g_a(\boldsymbol{x})},$$
(3.1)

where:

- \mathcal{F}_{Ω_t} is the 2D or 3D Fourier Transform for the *t*-th frame, associated with the sampling pattern Ω_t ;
- S_{ℓ} is the sensitivity map for the ℓ -th coil;
- $y_{t,\ell}$ is the *t*-th measurement on the ℓ -th coil;
- *T* is a sparsifying transform (*e.g.* Fourier transform, Finite Difference (for 1D-TV),
 ...) for the time dimension for each spatial location (voxel);
- $\Psi : \mathbb{C}^K \to \mathbb{C}^K$ is a sparsifying (*e.g.* wavelets, curvelets, etc.) transform along the spatial dimensions. The wavelet coefficient can be represented as $(\xi_a, (\xi_{s,j})_{s,1 \leq j \leq j_{max}})$ for each frame, where *s* is the orientation index (e.g., horizontal, vertical and diagonal in the case of dyadic decompositions) and *j* is the scale (*i.e.* resolution) index;
- $\underline{\Psi} = \Psi \otimes I_{N_t}$ is the Kronecker product of Ψ with the identity matrix of size N_t . This transform is applied to each frame independently;
- P_s extracts the detail wavelet coefficients $(\xi_{s,j})_{s,1 \leq j \leq j_{max}}$ for each frame t as the sparsity assumption is only tenable for these coefficients;
- P_a extracts the low-resolution approximation wavelet coefficients ξ_a for each frame;
- $\lambda_t, \lambda_s, \lambda_a > 0$ are the regularization parameters.

Remark. Typically, the data at hand would have a size $N_t \simeq 400, K = 192 \times 192 \times 128 \simeq 5 \cdot 10^6$ and $M \simeq \frac{10^6}{frame}$. The number of coils is usually L = 32, while the number of scales j_{max} depends on the spatial resolution, which in the fMRI context rarely exceeds 1 mm isotropic; therefore we consider $j_{max} = 4$.

It would

The remainder of this chapter consists of framing existing reconstruction methods within (3.1), as well as describing its extensions. The formulation of this problem is generic enough to encapsulate the classic approach for the reconstruction 3D + t data. In particular: If $\lambda_t = 0$, $\lambda_a = 0$ the problem is separable in time and each time frame can be reconstructed independently, this problem will be presented in § 3.1.2. Otherwise, the time dimension is considered, and this will be the focus of § 3.1.3.

3.1.2 Frame-wise reconstruction

Here, each frame is reconstructed independently, as follows (cf. Eq. (2.10)):

$$\boldsymbol{x}_{t} = \arg\min_{\boldsymbol{x}} \frac{1}{2} \sum_{\ell=1}^{L} \| \mathcal{F}_{\Omega_{t}} \mathcal{S}_{\ell} \boldsymbol{x} - \boldsymbol{y}_{t,\ell} \|_{2}^{2} + \mu g(\boldsymbol{\Psi} \boldsymbol{x})$$
(3.2)

for each frame index t.

Remark. This cost function is convex and of the form f + g, but usually non-differentiable (with $g = \ell_1$), and an optimal solution can be found iteratively using a wide range of algorithms such as ISTA (Theorem A.8) or accelerated version like FISTA or POGM (Theorem A.9). If g is differentiable (e.g. $g = \ell_2$), a simpler Conjugate Gradient Theorem A.7 could be used.

Warm start strategies

In practice, the image of frame t is not radically different from the data at time t - 1. Hence, to speed up the convergence of the optimization algorithm, the reconstruction can start from the last estimate: Indeed, since Eq. (3.2) is strongly convex, there exists a single minimizer to the problem, and initialization does not influence the result, only the number of iterations by positioning itself closer to the global minimizer.

However, in practice, the warm-start strategy can help boost the image quality, as will be shown in Chapter 7, especially if the trajectory Ω_t changes over consecutive frames.

Remark. A parallel can be drawn between the warm-start strategy and the stochastic gradient descent: Considering that \boldsymbol{x} does not evolve over time, we get a highly overdetermined set of observations $(\boldsymbol{y}_t)_{t \in [1,N_t]}$, and processing them all together would be too costly in practice. Instead, a batch of size 1 (1 frame) is selected and optimized. Repeating this many times will converge to the average estimate for all observations. More details on the convergence guarantees are available in [Kon+16; Nit14]. The application of these algorithms to the fMRI problem is one of the future perspective stemming form this thesis.

3.1.3 Time-aware reconstruction

The easiest image reconstruction scenario consists of setting $\lambda_a = \lambda_s = 0$ in Eq. (3.1), which is typically used in the context of dynamic MRI, with $\mathcal{T} = \mathcal{F}_T$, the temporal Fourier transform; this corresponds to the so-called *k-t FOCUSS* method [Jun+09]:

$$\hat{\boldsymbol{x}} = \operatorname*{arg\,min}_{\boldsymbol{x} \in \mathbb{C}^{N_t \times K}} f(\boldsymbol{x}) + \lambda_t \sum_{k=1}^K \|\mathcal{F}_T \boldsymbol{x}_{\cdot,k}\|_1.$$
(3.3)

However, stronger low-rank priors in fMRI are more effective as the brain is only subject to rigid motion, in contrast to the heart. This has thus led to the following methods.

Low-rank approaches. When $\lambda_s = \lambda_t = 0$ and $\Psi = Id$ in this case, we have the following problem:

$$\hat{\boldsymbol{x}} = \operatorname*{arg\,min}_{\boldsymbol{x} \in \mathbb{C}^{N_t \times K}} f(\boldsymbol{x}) + \lambda_a \|\boldsymbol{x}\|_*$$
(3.4)

This has been widely used in the k-t methods [Chi+16; OCS15; LF19; Chi+15; PHA17]

The low-rank + sparse case arises when $\lambda_s = 0$ and $\Psi = Id$ in this case, we end up with the following formulation:

Chapter 3 Image reconstruction methods for fMRI

$$\hat{oldsymbol{x}} = rgmin_{oldsymbol{x} \in \mathbb{C}^{N_t imes K}} f(oldsymbol{x}) + \lambda_t \sum_{k=1}^K \|\mathcal{T}oldsymbol{x}_{\cdot,k}\|_1 + \lambda_a \|oldsymbol{x}\|_*$$

The proximity operator of the nuclear norm is the thresholded singular Value decomposition.

Temporal Fourier sparsity. When \mathcal{T} is the temporal Fourier transform (applied to each voxel independently), the problem has been studied in the literature in two forms, the constrained and penalized ones.

Constrained Formulation. The problem can be formulated as a constrained optimization and solved using ADMM:

$$\hat{\boldsymbol{x}} = \hat{\boldsymbol{x}}_L + \hat{\boldsymbol{x}}_S \quad \text{where} \quad \begin{cases} (\hat{\boldsymbol{x}}_L, \hat{\boldsymbol{x}}_S) = \underset{\boldsymbol{x}_L, \boldsymbol{x}_S}{\arg\min \lambda_L \|\boldsymbol{x}_L\|_*} + \lambda_S \|\mathcal{T}\boldsymbol{x}_S\|_1 \\ \text{s.t.} \quad \mathcal{F}(\boldsymbol{x}_L + \boldsymbol{x}_S) = \boldsymbol{y} \end{cases}$$
(3.5)

where \boldsymbol{x}_L and \boldsymbol{x}_S denote respectively the low-rank and temporally sparse components of the 4D fMRI sequence, \mathcal{F} is the spatial frame-wise, Fourier transform, \mathcal{T} is the temporal sparsifying basis (*e.g.* the temporal Fourier in 1D transform \mathcal{F}_1) and \boldsymbol{y} is the measured data.

Penalized Formulation. The data consistency can also be expressed as:

$$(\hat{\boldsymbol{x}}_L, \hat{\boldsymbol{x}}_S) = \underset{\boldsymbol{x}_L, \boldsymbol{x}_S}{\operatorname{arg\,min}} f(\boldsymbol{x}_L + \boldsymbol{x}_S) + \lambda_L \|\boldsymbol{x}_L\|_* + \lambda_S \|\mathcal{T}\boldsymbol{x}_S\|_1.$$
(3.6)

Here we have $\lambda_L \equiv \lambda_a$ and $\lambda_S \equiv \lambda_t$ in Eq. (3.5) This approach has mainly been studied in cardiac MRI and fMRI, using a vectorized accelerated forward-backward algorithm. However, tuning the hyperparameters remains challenging.

Equivalence between the two formulations. Considering the augmented Lagrangian of Eq. (3.5), we obtain:

$$\mathcal{L}_{\rho}(\boldsymbol{x}_{L},\boldsymbol{x}_{S},\boldsymbol{\mu}) = \lambda_{L} \|\boldsymbol{x}_{L}\|_{*} + \lambda_{S} \|\mathcal{T}\boldsymbol{x}_{S}\|_{1} + \frac{\rho}{2} \|\mathcal{F}(\boldsymbol{x}_{L}+\boldsymbol{x}_{S})-\boldsymbol{y}\|_{2}^{2} + \langle \boldsymbol{\mu} | \mathcal{F}(\boldsymbol{x}_{L}+\boldsymbol{x}_{S})-\boldsymbol{y} \rangle$$
(3.7)

Then , setting $\tilde{\mu} = \mu / \rho$ and regrouping terms after scalar product expansion, we end up with:

$$\mathcal{L}_{\rho}(\boldsymbol{x}_{L},\boldsymbol{x}_{S},\tilde{\boldsymbol{\mu}}) = \lambda_{L} \|\boldsymbol{x}_{L}\|_{*} + \lambda_{S} \|\mathcal{T}\boldsymbol{x}_{S}\|_{1} + \frac{\rho}{2} \|\tilde{\boldsymbol{\mu}} + \mathcal{F}(\boldsymbol{x}_{L} + \boldsymbol{x}_{S}) - \boldsymbol{y}\|_{2}^{2} + \|\tilde{\boldsymbol{\mu}}\|_{2}$$
(3.8)

Eq. (3.8) is only equivalent to Eq. (3.6) if $\boldsymbol{\mu} = 0$ or $\boldsymbol{\mu} \perp \mathcal{F}(\boldsymbol{x}_L + \boldsymbol{x}_S - \boldsymbol{y})$.

If one can build a dictionary of time series (akin to the GLM in fMRI analysis), we can use it to regularize the temporal dimension of the data [CGM18] This can be formulated as:

$$(\hat{\boldsymbol{x}}_L, \hat{\boldsymbol{u}}) = \operatorname*{arg\,min}_{\boldsymbol{x}_L, \boldsymbol{u}} f(\boldsymbol{x}_L + \boldsymbol{u}\boldsymbol{v}^\top) + \lambda_L \|\boldsymbol{x}_L\|_*$$
(3.9)

where \boldsymbol{v} is the dictionary of time series, and \boldsymbol{u} is the loading vector for each time series in the dictionary. To ensure convergence, the constraint $\boldsymbol{x}_L \perp \boldsymbol{u} \boldsymbol{v}^{\top}$ is enforced by Graham-Schmidt orthogonalization.

The choice of the low-rank parameter is an implicit constraint on the rank of the matrix.

[CGM18] propose to control the contribution of the regularization terms with two hyperparameters τ, r , where r is the target rank of the matrix. Then the low-rank parameter is set to $\lambda_L = \tau \sigma_{r+1}(\boldsymbol{x})$, where $\sigma_r(\boldsymbol{x})$ is the r + 1 -th largest singular value of \boldsymbol{x} , *i.e.* the first value that would be discarded by a hard thresholding operator (In [CGM18] $\tau = 0.1, r = 16$ with 2 additional components of the low-rank + subspace dictionary).

3.2 Further extensions

The problem formulation in Eq. (3.1) can be extended in several ways to provide more freedom in the tuning of the parameters or reformulated to use more efficient computational implementations.

First, by considering vectorized regularization parameters. (e.g. voxel dependent λ_t , frame dependent λ_s , etc.) Then, the total variation (TV) regularization (case when $\mathcal{T} = \nabla_t$) can be relaxed by splitting the TV. In this case, we will have g^1, g^2 , which are independent functions. Further, the regularization in time g is also separable in spatial dimension, which means that the proximal operator for each voxel can be run independently. Similarly, the in-space regularization h is separable in the temporal dimension.

3.2.1 Other explicit priors

Sadly, Eq. (3.1) does not encapsulate all the possible priors. For instance, [Mas+21] propose other formulations (adapted here to be consistent with the notation of this chapter), considering the right and left subspaces of \hat{x} supposed to be of (low) rank r:

Explicit rank prior (kt-FASTER).

$$(\hat{\boldsymbol{u}}, \hat{\boldsymbol{v}}) = \operatorname*{arg\,min}_{\boldsymbol{u}, \boldsymbol{v}} \|\mathcal{F}(\boldsymbol{u}\boldsymbol{v}^T) - \boldsymbol{y}\|_2^2 \quad \text{s.t.} \quad \operatorname{rank}(\boldsymbol{u}) = \operatorname{rank}(\boldsymbol{v}) \leqslant r$$
(3.10)

where $\hat{\boldsymbol{x}} = \hat{\boldsymbol{x}}^L = \hat{\boldsymbol{u}}\hat{\boldsymbol{v}}^T$ is the final estimate.

Matrix Factorization.

$$(\hat{\boldsymbol{u}}, \hat{\boldsymbol{v}}) = \underset{\boldsymbol{u}, \boldsymbol{v}}{\operatorname{arg\,min}} \| \mathcal{F}(\boldsymbol{u}\boldsymbol{v}^{\top}) - \boldsymbol{y} \|_{2}^{2} + \lambda^{L_{u}} \|\boldsymbol{u}\|_{2}^{2} + \lambda^{L_{v}} \|\boldsymbol{v}\|_{2}^{2}$$
(3.11)

The model assumes that the data matrix can be represented as a product of two low-rank matrices, modeling the space $\hat{\boldsymbol{u}}$ and time $\hat{\boldsymbol{v}}$ dimensions, and regularize them using a ℓ_2 norm (also known as Thikonov or ridge regression).

Attachment to low resolution. By solving Eq. (3.10) with low resolution data (i.e. by windowing the k-space data to the center), we can obtain \boldsymbol{u}_{prior} and \boldsymbol{v}_{prior} , which can be used as a prior for the spatial and temporal subspace for the high resolution data.

$$\hat{\boldsymbol{u}}, \hat{\boldsymbol{v}} = \arg\min_{\boldsymbol{u}, \boldsymbol{v}} \| \mathcal{F}(\boldsymbol{u}\boldsymbol{v}^{\top}) - \boldsymbol{y} \|_2^2 + \lambda_{L_u} \| \boldsymbol{u} - \boldsymbol{u}_{prior} \|_2^2 + \lambda_{L_v} \| \boldsymbol{v} - \boldsymbol{v}_{prior} \|_2^2$$
(3.12)

The formulation of kt-PSF[Lam+13] is obtained by fixing the temporal subspace (ie $\boldsymbol{v} = \boldsymbol{v}_{prior}$) and only optimizes the spatial subspace.

$$\hat{\boldsymbol{u}} = \arg\min_{\boldsymbol{u}} \|\mathcal{F}(\boldsymbol{u}\boldsymbol{v}_{prior}^{\top}) - \boldsymbol{y}\|_{2}^{2} + \phi(\boldsymbol{u})$$
(3.13)

where $\phi(\boldsymbol{u})$ is a regularization function over the spatial subspace (ℓ_1 is used in [Lam+13]).

Temporal subspace smoothness. Given the assumption that the BOLD response maintains temporal smoothness, this prior can be enforced in the temporal dimension. This prior knowledge is generally expressed via TV regularization applied to the temporal subspace.

$$(\hat{\boldsymbol{u}}, \hat{\boldsymbol{v}}) = \arg\min_{\boldsymbol{u}, \boldsymbol{v}} \|\mathcal{F}(\boldsymbol{u}\boldsymbol{v}^{\top}) - \boldsymbol{y}\|_{2}^{2} + \lambda_{L_{\nabla}} \|\nabla_{t}\boldsymbol{v}\|_{2}^{2} + \lambda_{L_{u}} \|\boldsymbol{u} - \boldsymbol{u}_{prior}\|_{2}^{2} + \lambda_{L_{v}} \|\boldsymbol{v} - \boldsymbol{v}_{prior}\|_{2}^{2}.$$
(3.14)

Here, the TV is replaced by a norm ℓ_2 , giving an analytical solution. The algorithm alternatively minimizes the cost function (3.14) with respect to \boldsymbol{u} and \boldsymbol{v} .

3.3 Conclusion

This chapter explores recent progress in reconstructing fMRI image using compressed sensing techniques. There are numerous promising methods for reconstructing a sequence of volumes from 5D undersampled k-space data. Nonetheless, comparing these methods poses a challenge due to the lack of a definitive ground truth. To address this core issue, we will suggest a solution in Chapter 7. Additionally, the data consistency step emerges as the key bottleneck in reconstruction, requiring an efficient implementation of the MRI Fourier model. This aspect will be covered in Chapter 5. In Chapter 4 we will go one step further by considering the application of deep learning based method to the reconstruction of fMRI data.

Moreover, it is important to remind the reader that the reconstruction represents the initial stage of the processing pipeline. The quality of the reconstruction significantly influences the analysis that follows, which includes the image domain preprocessing pipeline and statistical testing. As illustrated in Chapter 6, further processing of the fMRI time series is essential for the elimination of thermal noise.

Die Gedanken sind frei, wer kann sie erraten, sie fliehen vorbei wie nächtliche Schatten. Kein Mensch kann sie wissen, kein Jäger sie schießen es bleibet dabei: Die Gedanken sind frei!

— German Volkslied

Chapter 4

Deep Learning for MRI and functional MRI

Outline

4.1	Deep Learning for inverse problems	57
	4.1.1 Supervised Learning	58
	4.1.2 Unsupervised and self-supervised learning	58
4.2	Neural Network architecture for inverse problems	59
	4.2.1 U-NET Network as a building block	59
	4.2.2 Algorithm unrolling	30
	4.2.3 Plug-and-Play approaches	31
4.3	Application to MRI And functional MRI	32
	4.3.1 Challenges for Deep Learning in MRI	32
	4.3.2 Deep Learning methods for fMRI reconstruction is even more	
	$challenging \dots \dots$	33
	4.3.3 An hybrid solution with Plug-and-Play priors	33
4.4	Conclusion	34

C HAPTER 3 presented variational reconstruction methods that use convex optimization to solve an inverse problem. Yet, in the last decade, with the joint increase in computational power and data availability, deep learning methods have emerged as a promising solution to solve inverse problems¹ of the form $y = Ax + \epsilon$, which is the case for MRI and fMRI reconstruction. In the following, we will first introduce the general concepts and methods used to solve inverse problem using deep learning, then we will present their application to the MRI reconstruction problem and discuss the existing challenges and limitation for their broader adoption and extension to fMRI reconstruction.

4.1 Deep Learning for inverse problems

So far, we have shown how we could solve an inverse problem by defining a convex objective function and solving it using iterative algorithms. We have access to a wide variety of tools to solve these problems, carefully crafted prior and data-consistency terms, with a good understanding on how to use them.

¹The unfamiliar reader can refer to Chapter A for an introduction to inverse problems

In this section, we forgo all of this, and instead solve the problem by relying on large datasets and computing power. Namely, one can train a Deep Neural Network (DNN) G on a dataset of paired desired and measured values $(\boldsymbol{x}_n, \boldsymbol{y}_n)_{n \in \mathbb{N}}$ so that $G(\boldsymbol{y}) \simeq \boldsymbol{x}$, in a black-box fashion. This has been shown to be very effective, as long as a large collected of paired data (supervised setting) are available to uncover G using specific assumptions and a given architecture of deep neural network.

Remark. In the statistics and deep learning community the roles of \boldsymbol{x} and \boldsymbol{y} are swapped compared to the imaging community we used in this manuscript: \boldsymbol{x} represents the, measured, or observed, input data and \boldsymbol{y} the reference, or output label, ground-truth. Here, \boldsymbol{x} represents the reference or ground truth images whereas \boldsymbol{y} represents the undersampled k-space data.

Remark. This section is only a shallow overview of the field, and is not meant to be exhaustive. The reader is encouraged to refer to the following books and articles for a more in-depth understanding of the field and its application to inverse problems solving [Sca+22; AMJ19].

4.1.1 Supervised Learning

Let $\mathcal{H}^N, \mathcal{G}^N$ be two Hilbert spaces that represent respectively, the image domain and the measurement domain Given a DNN G, a training set $\mathcal{D} = (\boldsymbol{x}_n, \boldsymbol{y}_n)_{1,\dots N} \in \mathcal{H}^N \times \mathcal{G}^N$, and a loss function \mathcal{L} , training the network is done by minimizing the following objective:

$$\hat{\theta} = \arg\min_{\boldsymbol{\mu}} \mathbb{E}_{\boldsymbol{x}, \boldsymbol{y} \in \mathcal{D}} \mathcal{L}(G_{\theta}(\boldsymbol{y}_i), \boldsymbol{x}_i)$$
(4.1)

where θ denotes the learnable parameters of $G_{\theta} : \mathcal{G} \to \mathcal{H}$ and $\mathcal{L} : \mathcal{H} \to \mathbb{R}$ is a loss function.

Unlike the previous sections, (4.1) is typically not a convex objective, meaning that none of the earlier solver can guarantee an accurate estimation of θ . However, since G_{θ} has been built from simple functions, it is inherently differentiable. This allow us to leverage the gradient of the loss function to update the parameters., and we can use the gradient of the loss function to update the parameters θ .

Furthermore, the number of parameters in θ is typically very large (in the order of 10⁷), and the optimization problem is highly non-linear, so that the optimization landscape is very complex. Thus, the optimization problem is typically solved using a variant of the Stochastic Gradient Descent (SGD) algorithm, which is a variant of the gradient descent algorithm, where the gradient is computed on a subset of the data (a *batch*) and the step size can be adapted during the optimization to ensure convergence to a local minimum.

4.1.2 Unsupervised and self-supervised learning

All the aforementioned approaches requires access to a dataset of fully-sampled data for training. However, as discussed in Chapter 2 it is common practice in MRI to acquire under-sampled data. Training with under-sampled data has led to the development of so-called *self-supervised* techniques. In these approaches, the under-sampled data is

further divided into two subsets [Akc+21; Yam+20a], and the optimization objective typically takes the following form:

$$\arg\min_{\theta} \mathcal{L}(\boldsymbol{y}_{(2)}, \boldsymbol{A}_{(2)}(G_{\theta}(\boldsymbol{y}_{(1)})))$$
(4.2)

Where $\boldsymbol{y}_{(i)}$ is the split of the train /test data in the measurement, and $\boldsymbol{A}_{(2)}$ the operator associated with the test measurement. Various network architectures have been proposed to implement this optimization framework effectively, including Variational AutoEncoder (VAE), Generative Adverserial Networks (GAN)[QNJ18], and Visual Transformer (ViT)[LH22].

Remark. The self supervised learning can also be used for training denoising network [XWJ20; KBJ19; Hua+21; Leh+18], such network can then be used in plug-and-play framework (see § 4.2.3), or to pre-process the dataset before training the network on cleaner data (an approach proposed in Chapter 8).

4.2 Neural Network architecture for inverse problems

Using neural networks to solve inverse problems seems quite natural, especially when sufficient paired training samples of source and measured data are available. However, blindly attempting to learn the reverse mapping from the data to the source is not always the optimal approach. This so-called *physics- blind* approach has been shown to be quite limited in practice [Zhu+18],and does not scale well to high-dimensional data. Additionally, the network must expend significant effort to learn the forward model of the problem, which is typically already known and has an efficient implementation (e.g., the FFT for the Fourier transform). As a result, such networks are more sensitive to hallucination and overfitting.

4.2.1 U-NET Network as a building block

Apart from the physics model, the structure of the data itself can significantly influence the design of the network. For example, images contain spatial information that can be effectively exploited by Convolutional Neural Networks (CNNs), while time-series data is better suited for processing by Recurrent Neural Networks (RNNs) or transformers. Additionally, the varying physical scales of the data can be optimally handled when the receptive field of the network is adapted to match these scales. To address this, networks can be designed with multi-scale architectures, such as the U-Net [RFB15], illustrated in Figure 4.1.

U-Net architectures have shown good performance well in image denoising tasks² (the simplest inverse problem where $\mathbf{A} = \mathbf{I}$) [Zha+21b]. In essence, the U-Net architecture consists of a symmetric encoder-decoder structure. The encoder (downsampling path) reduces the spatial dimensions while increasing the number of feature channels through convolution and pooling, thereby capturing high-level semantic features. The decoder (upsampling path) restores spatial resolution using transposed convolutions to reconstruct finer details.

²But not only! Segmentation tasks are also a well-known application



Figure 4.1: U-Net Architecture. Adapted from [Zha+21b]

A key feature of the U-Net is the use of skip connections, which directly link corresponding stages of the encoder and decoder. These connections help preserve spatial details lost during downsampling and address challenges such as the vanishing gradient problem, ensuring more robust and accurate reconstructions.

A straightforward approach is to use a U-Net following the physics model operator to improve the quality of reconstructed data and reduce undersampling artifacts. However, such networks lack an inherent connection to the original data, as the physics model is not embedded within their architecture. Consequently, data fidelity cannot be guaranteed, which may result in inaccurate or misleading outcomes.

In the following sections we will discuss the two methods that are most commonly used in practice: Algorithm unrolling (§ 4.2.2) and Plug-and-Play (§ 4.2.3), with their main architectures illustrated in Figure 4.2

4.2.2 Algorithm unrolling

To overcome the limitations of these physics-blind approaches, a natural strategy is to introduce a data consistency measure using the classic Gradient descent step $\boldsymbol{x} \rightarrow \boldsymbol{x} - \eta \nabla f(\boldsymbol{x})$ (see Theorem A.7), and then chain this data-fidelity term with the DNN architecture multiple time.

Effectively, we have recreated an iterative algorithm, alike the Forward-Backward algorithm, but with a DNN instead of the proximal operator. This approach is called *algorithm unrolling* and has shown to be very effective in practice [AO18; Ham+18].

For training such models, several design choices can be made: The weights of all the DNN can be shared (as proposed in MoDL [AMJ19]), or each DNN can have its own set of weights (as in the Deep Cascade [Sch+18]). This approach can be pushed further by putting DNN on both sides of the physics model, such as in the primal-dual network [AO18]. For MRI, this has been put in practice in the XPDNET and NCPDNET architectures [RCS21; Ram+22], in earlier work of our team.

Remark. While unrolled approaches offer significant benefits, they are not a panacea. They require large datasets and substantial computational resources for training. Additionally, they rely on the forward model being known and differentiable in order to train the model. In practice, this has limited their adoption in non-Cartesian MRI reconstruction, as efficient implementations were not readily available.

A typical use of unrolled networks is to learn how to reconstruct high-resolution data


Figure 4.2: Neural Network design for Inverse problem, and how they are trained in the context of MRI. *Top left*: End-to-end/Physics blind network are trained without embedding physic knowledge. *Top right*: Unrolled Networks trains on the reconstruction process. *Bottom*: The Plug-and-Play approach uses a separately trained network on a denoising task. Adapted from [Ter23].

from under-sampled measurements, with fully sampled data available in a supervised setting. Once trained, inference for data reconstruction is significantly faster than the algorithms discussed in the previous section. However, a drawback of this approach is that the network is typically trained for a specific acceleration factor (AF) and under-sampling pattern. Since the physics is deeply integrated into the network, it tends to generalize poorly to unseen acquisition scenarios at training, i.e., other acceleration factors and sampling patterns

4.2.3 Plug-and-Play approaches

By combining the strengths of both convex iterative optimization and deep learning, the Plug-and-Play (PnP) approach focuses on learning the prior g, specifically its proximal operator. The proximal operator of the prior can be viewed as a denoising function, and DNNs have proven to be highly effective in denoising tasks. The core idea is to replace the proximal operator with a DNN in one of the iterative algorithms presented in Chapter A, adapting the proximal gradient descent Theorem A.8 to its PnP version (4.3).

$$\forall k = 0, \dots, N \qquad \forall k = 0, \dots, N \\ \begin{bmatrix} \boldsymbol{z}_k = \boldsymbol{x}_k - \eta_k \nabla f(\boldsymbol{x}_k) & \xrightarrow{PnP} \\ \boldsymbol{x}_{k+1} = \boldsymbol{x}_k + \lambda_k (\operatorname{prox}_{\eta_k g}(\boldsymbol{z}_k) - \boldsymbol{x}_k) & \begin{bmatrix} \boldsymbol{z}_k = \boldsymbol{x}_k - \eta_k \nabla f(\boldsymbol{x}_k) & (4.3) \\ \boldsymbol{x}_{k+1} = G_{\theta}(\boldsymbol{z}_k) \end{bmatrix}$$

This approach has been shown to be highly effective in practice, and can be applied to a wide range of inverse problems, including denoising, deblurring, in-painting, and super-resolution [VBW13; CWE16].

Many methods along with their specific tuning can be understood as a Plug-and-Play approach when the prior g is learned: Examples include Deep Image Prior [UVL20], RED [REM17] and even Dictionary Learning.

Moreover, several parallels can be drawn between PnP and diffusion models used for generative imaging[Sun+24; Wu+24]. Specifically, guided diffusion model with data consistency are very similar to the PnP methods, while also offering the valuable ability to generate multiple samples from the posterior, which is instrumental to perform uncertainty quantification of the reconstructed data. However, such methods are typically more expensive to compute [Lau+22].

4.3 Application to MRI And functional MRI

The inverse problem formalism of the fMRI reconstruction also allows the use of deep learning-based methods. In § 4.1 we presents a general overview of the available methods for solving the general inverse problem.

The use of deep-learning based methods for MRI is an ever-growing field, and the reader is encouraged to refer to the recent review [Hec+24]. Let us at least briefly mention that most of the development in this field has been focused on anatomical MRI and generally restricted to 2D images, and have used supervised dataset such as fastMRI [Zbo+18] or Calgary-Campinas[Sou+18]. In these settings, it has been proven many times that deep learning can outperform traditional compressed sensing methods[Hec+24], especially when using physics-informed neural networks [AO18] and unrolled approaches [Ham+18; RCS21].

However, extending these methods to fMRI presents several challenges that needs to be addressed.

4.3.1 Challenges for Deep Learning in MRI

The successful deployment of deep learning-based model relies heavily on various assumptions, which may not always hold true, even for anatomical MRI. This section highlights several limitations and challenges that need to be addressed in training these complex systems.

- *Overfitting and hallucinations*: One major concern is the tendency of Deep Neural Network to overfit to the training data, leading to hallucinations at test time.
- *Dimensionnality of the dataset*: Scaling up to more complex data types with higher dimensions, such as 3D or multichannel data, poses significant challenges due to

the large literature already established for 2D natural images. Adding a temporal dimension further complicates training, by introducing memory bottlenecks, and extra computations.

- *Large dataset availability*: The availability of large datasets remains a prerequisite for successful training of ever bigger and more powerful networks.
- Data crimes: When processing the dataset, a careful methodology needs to be established, aside from the –unfortunately frequent– train/test split leaks. In particular preprocessing or limitation of the dataset (typically using magnitude-only images in MRI context) introduces bias in the data, that would be transferred in the networks weights [Shi+22].
- Generalization to other dataset and sampling scheme: A Deep Neural Network trained on a specific dataset will embeds in its weights implicit prior reflecting the potential bias of this data, in particular, unrolled network generalizes poorly to new data acquired from different MRI hardware, or with new acquisition sampling patterns.

4.3.2 Deep Learning methods for fMRI reconstruction is even more challenging

Applying Deep learning to fMRI reconstruction is even more complex than for anat MRI. One significant challenge in implementing deep learning techniques for fMRI data is the lack of fully sampled high-resolution datasets, which limits the possibility of supervised training with corresponding ground truth. The trade-off between spatial and temporal resolution during acquisition (as discussed in Chapter 1 and Chapter 2) imposes a physical constraint to the acquisition of such data. Most available fMRI datasets are designed for specific processing tasks and do not provide direct access to raw k-space data, hindering the development of more advanced temporal models. Moreover, many deep learning challenges in MRI become more pronounced along the temporal dimension, requiring substantial memory and computational resources for training. This scalability issue remains a key barrier to unlocking the full potential of deep learning methodologies in fMRI image reconstruction and analysis.

4.3.3 An hybrid solution with Plug-and-Play priors

In § 4.2.3, we introduced the Plug-and-Play framework, where the explicit regularization prior is replaced by a learned denoiser acting as a proximal operator.

In this scenario, regularization functions—possibly several—can be substituted with a learned prior. An initial examination of their implementation for magnetic resonance imaging is presented in Chapter 8, along with contributions aimed at accelerating the convergence of this iterative scheme using pre-conditioning. Importantly, these approaches offer an effective balance between computational demands and efficiency compared to fully unrolled methods, and their application is flexible across different acquisition settings, as they are agnostic to the forward operator.

Extending the Plug-and-Play framework to fMRI reconstruction is an on-going effort at the time of writing this manuscript. In particular, the temporal dimension of the data can be used for training a denoiser in a self-supervised setting (akin to the SSDU framework[Yam+20a; Yam+20b; Akç+21], or using diffusion models [Gao+25]).

4.4 Conclusion

In this chapter, we outlined the main concepts and methods used to solve inverse problems with deep learning and discussed their application to MRI and fMRI reconstruction. However, while using deep learning for MRI is already challenging, extending it to functional MRI exacerbates the problem. Specifically, ensuring the fidelity of the reconstruction is crucial, and guarantees for downstream analysis should ideally be provided. In other words, the reconstruction should avoid hallucinating additional activations, while preserving both the spatial and temporal structure of the data. Among the various solutions proposed in the literature for MRI (such as algorithm unrolling and physics-informed networks), the Plug-and-Play framework presents a promising approach that could be scaled to address these challenges and is currently being explored for fMRI reconstruction. A preliminary attempt in this is provided for anatomical MRI in Chapter 8.



Summary of concerns

Non-Cartesian fMRI is computationally demanding.

While Non-Cartesian MRI allows for faster acquisition of under-sampled data, it requires a dedicated reconstruction setup, typically involving a Non-Uniform Fast Fourier Transform (NUFFT), an extended forward model (parallel imaging, field inhomogeneities, etc.) and an iterative optimization algorithm. Functional MRI intensifies this burden by introducing the extra temporal dimension.

To efficiently explore new fMRI acquisition and reconstruction methods, it is essential to have access to high-performance NUFFT computation with efficient memory management. Chapter 5 presents our proposal to address this challenge.

How do you measure the performance of a fMRI reconstruction ?

Measuring the performance of fMRI reconstruction is challenging due to the reliance on BOLD contrast as a proxy for brain activity, which means we cannot directly access the true activation state. Consequently, it is impossible to guarantee that the reconstructed data faithfully capture all neuronal activity locations. Additionally, factors such as subject variability and noise sources limit the *in-vivo* reproducibility of fMRI studies.

Before advancing such strategies, it is essential to establish reliable tools to evaluate the quality of reconstructed images, using high-quality fully sampled reference k-space data. Chapter 6 offers a preliminary attempt to denoise experimental data for this purpose. However, Chapter 7 presents our primary contribution, providing a comprehensive approach to designing an artificial fMRI data simulator to validate the data acquisition and image reconstruction up to the statistical analysis.

Could Deep-Learning methods help reconstructing fMRI data ?

Deep learning reconstruction methods that are effective for anatomical MRI cannot be directly applied to functional MRI (fMRI) without significant modifications. Specifically, supervised training of unrolled networks is impractical for fMRI due to the lack of fully sampled raw fMRI datasets. Moreover, creating such datasets would require compromising either the spatial or temporal resolution. The computational and memory demands for training these networks are also a challenge, given the limitations of current hardware. Additionally, changes in acquisition strategies would necessitate a new and costly training process for the unrolled network.

To address these challenges, more efficient methods that can generalize across a variety of acquisition settings and reconstruction tasks are essential to avoid unnecessary costs in fMRI research. In Chapter 8, we explore a potential solution in the form of the Plug-and-Play Network, which offers a more flexible and scalable approach for fMRI reconstruction.

Part III

Methodological developments

Try and leave this world a little better than you found it. — Robert Baden Powell

Chapter 5

MRI-NUFFT: Doing non-Cartesian MRI has never been easier.

Outline

Con	text and summary
5.1	Features
	5.1.1 NUFFT Library compatibility
	5.1.2 Extended Fourier Model
	5.1.3 Density Compensation
	5.1.4 Trajectories generation and expansions
	5.1.5 Auto-differentiation for data and sampling pattern
5.2	Benchmark of existing NUFFT libraries
5.3	Example of Non-Cartesian, Non-Fourier Reconstruction using MRI-
	NUFFT and compressed sensing
	5.3.1 Material and methods \ldots \ldots \ldots \ldots \ldots $.74$
	5.3.2 Results \ldots \ldots $.$ 76
5.4	Conclusion

The following is a combination of two collaborative efforts:

- Pierre-Antoine Comby, Guillaume Daval-Frérot, Caini Pan, Asma Tanabene, Léna Oudjman, Matteo Cencini, Philippe Ciuciu, and Chaithya Gr. MRI-NUFFT: Doing Non-Cartesian MRI Has Never Been Easier. Version 2. Nov. 2024. URL: https://hal.science/hal-04775262 (visited on 12/12/2024). in revision to Journal of Open Source Software.
- Z. Amor, P.-A. Comby, C. Le Ster, A. Vignaud, and P. Ciuciu. "Non-Cartesian Non-Fourier FMRI Imaging for High-Resolution Retinotopic Mapping at 7 Tesla". In: 2023 IEEE 9th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP). Dec. 2023, pp. 201–205
- P.-A. Comby, G. Daval-Frerot, Chaithya Gr, A. Vignaud, and P. Ciuciu. "MRI-NUFFT: An Open Source Python Package to Make Non-Cartesian MR Imaging Easier". In: *ISMRM Annual Meeting, (in Press)*. Singapore, 2024

Context and summary

As presented in Chapter 1 and Chapter 2, MRI is a non-invasive biomedical imaging technique, where raw data is sampled in the spatial frequency domain (k-space) and final images are obtained by applying an inverse (fast) Fourier transform on this data. Traditionally, the data is sampled on a Cartesian grid (often partially by skipping lines to accelerate the acquisition) and reconstructed using FFT-based algorithms. However, the Cartesian approach is not always the best choice for data collection, and non-Cartesian sampling schemes have been proposed to improve image quality, reduce acquisition time or enable new imaging modalities. Nonetheless, the reconstruction of non-Cartesian data is more challenging and requires the use of non-uniform fast Fourier transform (NUFFT) algorithms.

Several NUFFT libraries have been developed in the past few years, but they are not always easy to use or don't account for the specificities of MRI data acquisition (*e.g.* multi-coil acquisition, static-field inhomogeneities, density compensation, etc.). Also, their performances can vary a lot depending on the use cases (2D vs 3D data, number of coils, etc.).

To ease the use of NUFFT in the MRI (and fMRI) context, we developed MRI-NUFFT an open-source Python library that provides a universal interface to the different existing NUFFT library (CPU and GPU based), enriching them with an extended forward model for MRI. It also provides a wide collection of non-Cartesian sampling patterns that one can easily experiment with.

Furthermore, there has been a growing interest in using deep learning to jointly learn MRI acquisition and reconstruction, which requires computing the gradients of the loss function for image reconstruction with respect to the raw data and/or the sampling locations.

Finally, let us note that MRI-NUFFT is the underlying library that we are going to use in SNAKE (presented in Chapter 7) to perform the generation of k-space data.

5.1 Features

5.1.1 NUFFT Library compatibility

MRI-NUFFT is compatible with the following NUFFT librairies: FINUFFT [BMK19], CUFINUFFT [Shi+21], gpuNUFFT [Kno+14], TorchKbNufft [Muc+20], pyNFFT, SigPy [OL19] and BART [UOT15]. Using our benchmark we can also determine which NUFFT implementation provides the best performances both in term of computation time and memory footprint. At the time of writing, cufinufft and gpunufft provide the best performances by leveraging CUDA acceleration. MRI-NUFFT supports as well standard array libraries (NumPy, CuPy, PyTorch, TensorFlow, etc.) and optimizes data copies, relying on the array-API standard. It also provides several enhancements on top of these backends, notably an optimized 2.5D NUFFT (for stacks of 2D non-uniform trajectories, commonly used in MRI), and a data-consistency term for iterative reconstruction ($\mathcal{F}^*_{\Omega}(\mathcal{F}_{\Omega}\boldsymbol{x} - \boldsymbol{y})$).



Figure 5.1: MRI-NUFFT as an interface for non-Cartesian MRI

5.1.2 Extended Fourier Model

MRI-NUFFT provides a physical model of the MRI acquisition process, including multicoil acquisition and static-field inhomogeneities (see (1.17)). This model is compatible with the NUFFT libraries, and can be used to simulate the acquisition of MRI data, or to reconstruct data from a given set of measurements. Namely, we provide a linear operator that encapsulates the forward and adjoint NUFFT operators, the coil sensitivity maps and (optionally) the static field inhomogeneities. The forward model is described by the following equation:

$$y(\boldsymbol{\nu}_i) = \sum_{j=1}^N x(\boldsymbol{r}_j) e^{-2i\pi \boldsymbol{r}_j \cdot \boldsymbol{\nu}_i} + n_i, \quad i = 1, \dots, M$$
(5.1)

where: $x(\mathbf{r})$ is the spatially varying image contrast acquired; y_1, \ldots, y_M are the sampled points at frequency locations; $\Omega = \{\boldsymbol{\nu}_1, \ldots, \boldsymbol{\nu}_M \in [-1/2, 1/2]^d\}$; \mathbf{r}_j are the N spatial locations of image voxels, and n_i is a zero-mean complex-valued Gaussian noise, modeling the thermal noise of the scanner.

This can also be formulated using the operator notation

$$\boldsymbol{y} = \mathcal{F}_{\Omega}(\boldsymbol{x}) + \boldsymbol{n} \tag{5.2}$$

As the sampling locations Ω are non-uniform and the image locations u_j are uniform, \mathcal{F}_{Ω} is a Non Uniform Discrete Fourier Transform (NUDFT) operator, and the equation above describe a Type 2 NUDFT. Similarly, the adjoint operator is a Type 1 NUDFT, see Table 5.2.

NUFFT Type	Operation	MRI Transform	Operator
Type 1 Type 2	Adjoint Forward	$\begin{array}{l} \text{K-space} \rightarrow \text{image} \\ \text{Image} \rightarrow \text{k-space} \end{array}$	${\mathcal F}^*_\Omega \ {\mathcal F}_\Omega$

 Table 5.2: Equivalence between NUFFT operations

Parallel Imaging Model

As explained in Chapters 1 and 2, the acquired signal can be received by multiple antennas ("coils"). Each coil possesses a specific sensitivity profile (i.e. each sees the object differently due to its physical layout).

$$\tilde{\boldsymbol{y}} = \begin{bmatrix} \mathcal{F}_{\Omega} S_1 \\ \vdots \\ \mathcal{F}_{\Omega} S_L \end{bmatrix} \boldsymbol{x} + \boldsymbol{n}_{\ell} = \mathcal{F}_{\Omega} S \otimes \boldsymbol{x} + \tilde{\boldsymbol{n}}, \qquad (5.3)$$

where S_1, \ldots, S_L are the sensitivity maps of each coil. Such maps can be acquired separately by sampling the k-space low frequencies, or estimated from the data.

Off-resonance correction model

Eq. (1.15) assumes a homogeneous B_0 and highly precise application of the prescribed trajectories ($\mathbf{k} = [\mathbf{k}_s]_{s \in [1, N_{shot}]}$) by the system. Unfortunately, in the case of highly undersampled trajectory (for instance 3D-SPARKLING[Cha+22]), a strong degradation in image quality occurs due to B_0 imperfections arising from the patient (susceptibility changes at the tissue-air interfaces, physiological noise) and the system (eddy currents, temperature drifts, system instabilities) [Kas+15; Van+13; Bol+17].

One way to overcome this issue is to take into account such imperfections in the signal model during reconstruction: Eq. (1.15) then becomes Eq. (5.4) where $\Delta B_0^{\text{stat}}(\mathbf{r})$ (in Hz) denotes the static inhomogeneities of the B_0 field and $\Delta B_{0,s}^{\text{dyn}}$ (in Hz) and $\mathbf{k}_s = \mathbf{k}_s + \delta \mathbf{k}_s$ (in m⁻¹) denote respectively its zeroth order dynamic fluctuation and the measured trajectory (deviated from the prescribed one due to first order fluctuation $\delta \mathbf{k}_s$). $\Delta B_{0,s}^{\text{dyn}}$ is slowly varying and considered constant during a shot.

$$y_{\ell}^{s}(t) = e^{-2i\pi t \Delta B_{0,s}^{\text{dyn}}} \int_{\text{FOV}} \bar{x}_{\ell}(\boldsymbol{r}) e^{-2i\pi \left[\Delta B_{0}^{\text{stat}}(\boldsymbol{r})t + \tilde{\boldsymbol{k}}_{s}(t) \cdot \boldsymbol{r}\right]} \,\mathrm{d}\boldsymbol{r} \,.$$
(5.4)

According to (5.4), the discretized backward operator can be written as in (5.5):

$$\bar{x}_{\ell}(\boldsymbol{r_n}) = \sum_{t} e^{2i\pi t \Delta B_{0,s}^{\text{dyn}}} y_{\ell}^s(t) e^{2\pi i \left[\Delta B_0^{\text{stat}}(\boldsymbol{r})t + \tilde{\boldsymbol{k}}_s(t) \cdot \boldsymbol{r}\right]}.$$
(5.5)

The mixed term $e^{2i\pi\Delta B_0^{\text{stat}}(\mathbf{r}_m)t} = \sum_P b_{m,p}c_{p,n}$ is split in a rank-*P* linear decomposition using a SVD [Fes+05; Dav+22] and we end up with a sum of *P* (non-uniform) Fourier transform in (5.6).

$$\bar{x}_{\ell}(\boldsymbol{r}_{n}) = \sum_{p=1}^{P} c_{p,n} \sum_{t_{m}=0}^{T_{obs}} b_{m,p} \underbrace{e^{2i\pi t \Delta B_{0,s}^{dyn}} y_{\ell}^{s}(t_{m})}_{\tilde{y}_{\ell}^{s}(t_{m})} e^{2i\pi \tilde{k}_{s}(t_{m}) \cdot \boldsymbol{r}_{n}} .$$
(5.6)

Since the term related to $\Delta B_{0,s}^{\text{dyn}}$ is outside of the integral in Eq. (5.4), the zerothorder dynamic fluctuations can be corrected simply by demodulating each $y_{\ell}^{s}(t)$ by the corresponding $e^{2i\pi t\Delta B_{0,s}^{\text{dyn}}}$ to obtain $\tilde{y}_{\ell}^{s}(t)$.

Eq. (5.6) holds for all frequencies $\mathbf{k}(t_m)$ and locations \mathbf{r}_n across all N_{shot} we can summarize the perturbed acquisition in Eq. (5.7), as a sum of the adjoint non-uniform Fourier transform $\mathcal{F}_{\tilde{\Omega}}$, which produces a coil image \mathbf{x}_{ℓ} from the measured frequency locations $\tilde{\Omega}$ and associated corrected values $\tilde{\mathbf{y}}_{\ell}$.

$$\bar{\boldsymbol{x}}_{\ell} = \sum_{p=1}^{P} \boldsymbol{c}_{p} \odot \boldsymbol{\mathcal{F}}_{\Omega}^{*}(\boldsymbol{b}_{p} \odot \tilde{\boldsymbol{y}}_{\ell}) = \tilde{\boldsymbol{\mathcal{F}}}^{*}(\tilde{\boldsymbol{y}}_{\ell})$$
(5.7)

where \odot denotes the element-wise product. Using this extended Fourier operator and at the cost of additional computations, we can adjust the cost function in (2.10) for the CS reconstruction, applied to each volume independently.

5.1.3 Density Compensation

In non-Cartesian acquisition, the sampling frequencies locations are typically not uniformly distributed. For instance, there are more points sampled in the center of k-space than in its periphery for radial or spiral trajectories. Thus, the lower frequencies are over-emphasized in the adjoint $\mathcal{F}^*_{\Omega} \boldsymbol{y}$ and leads to blurred images and bad conditioning of the forward model operator (see Definitions A.1 for more details), even if the sampling is done at Nyquist's rate. To correct for this effect, density compensation was introduced.

The main idea is to reweight the k-space values with a diagonal matrix \boldsymbol{W} , obtained from a given heuristic. The end goal is for the operator $\mathcal{F}^*_{\Omega} \boldsymbol{W} \mathcal{F}_{\Omega}$ to be close to identity.

Determining a perfect W has been an active area of research, with the main idea of getting a fast estimation of W. Typical choice are 1. Voronoi Parcellation; 2. Cell Couting; 3. Pipe's Iterative scheme; 4. Analytic solutions (only viable for light parametrized trajectory, like radial).



Figure 5.3: Effect of density compensation for a radial trajectory. Applying the density compensation (here) Voronoi restores the contrast of the image, but also introduces some high frequency artifacts.

Remark. Density compensation is not a perfect solution. Notably, as it changes the definition of the forward operator, and put more emphasis on high frequencies than the original operator, it amplifies noise. Furthermore, similar image quality can be achieved by using an iterative scheme such as a Conjugate Gradient Descent (Theorem A.7). Nonetheless, a density-compensated adjoint image can be used a good initial solution for such iterative reconstruction (or ones using Compressed Sensing regularization like in Chapter 2 and Chapter 3). Density compensation has also shown to help training Deep Unrolled Network (§ 4.1) as it helps maintain the norm of images through the network.

5.1.4 Trajectories generation and expansions

MRI-NUFFT comes with a wide variety of non-Cartesian trajectory generation routines that have been gathered from the literature. It also provides ways to extend existing trajectories and export them to specific formats, for use in other toolkits and on MRI hardware.

5.1.5 Auto-differentiation for data and sampling pattern

Following the formulation of [WF23], MRI-NUFFT provides automatic differentiation for all NUFFT backends, with respect to both gradients and data (image or k-space). This enables efficient backpropagation through NUFFT operators and supports research on learned sampling model and image reconstruction network.

5.2 Benchmark of existing NUFFT libraries

In Figure 5.4 we report a benchmark on both computational time and memory demand of different NUFFT libraries available, in the context of MRI. We tested forward , adjoint, and data-consistency steps that are available in MRI-NUFFT with increasing number of coils. We report overall computational time and memory consumption.

From these experiments, it is clear that GPU-based implementation such as gpuNUFFT and cufinufft provides the best results, with cufinufft proving to be the fastest in 2D and in some cases of 3D trajectories. gpuNUFFT however remains more memory efficient, as it only instantiate a single *NUFFT Plan* for both forward and adjoint operation.

5.3 Example of Non-Cartesian, Non-Fourier Reconstruction using MRI-NUFFT and compressed sensing

5.3.1 Material and methods

The SPARKLING sampling pattern has been developed to optimize the sampling density in the k-space while fulfilling the hardware constraints on the magnetic field gradient system [Cha+22]. Its application to fMRI is thus a natural extension of the work done

¹The full benchmark is available at https://github.com/mind-inria/mri-nufft-benchmark/



Figure 5.4: Benchmarking of high resolution NUFFT operations¹

in the anatomical case. However, as a non-Cartesian sampling pattern, its sensitivity to field inhomogeneities needs to be addressed directly in the reconstruction process.

In the first attempt, the 3D-SPARKLING pattern is repeated, and each TR-specific reconstruction is performed scan-wise. However, a naive application of the CS theory as that of Eq. (2.10) does not account for static and dynamic field inhomogeneities. For optimal performance, it is thus required to model the effect of field inhomogeneities and measure them using an external device.

Remark. In general, calling $\tilde{\mathcal{F}}$ requires the computation of $P \times L$ NUFFTs. Although strategies to reduce the computational cost can be elaborated (such as coil compression or limiting the number of interpolators), this formulation for CS MR image reconstruction and its original implementation using CPU-bounded computation (notably for the wavelet synthesis/analysis transforms) required the use of a dedicated cluster. More precisely, image reconstruction took 10 hours per run (each run consisting of 125 volumes) on a 4 GPU node on the Jean-Zay supercomputer².

5.3.2 Results

In Figure 5.5, we outline the benefit of taking the field inhomogeneities up to the first order $(\Delta B_0^{\text{stat}}, \Delta B_{0,s}^{\text{dyn}} \text{ and } \delta \mathbf{k})$ in the forward model into account. Additional results are available in [Amo+23].

5.4 Conclusion

MRI-NUFFT has become a fundamental tool for the development of non-Cartesian MRI. It provides an unified Forward model that can be used in both variational or deep-learning based methods. As such, it will also be used internally in the SNAKE simulator (Chapter 7), and in Chapter 8.

MRI-NUFFT also gives access to a collection of existing methods to create, augment and visualize non-Cartesian sampling patterns. More examples of its uses are available in its documentation³.

Finally, MRI-NUFFT has been successfully used in experimental settings for the reconstruction of high spatial resolution fMRI data, embedding an extended model correcting off-resonance effects.

²http://www.idris.fr/eng/jean-zay/

³https://mind-inria.github.io/mri-nufft



Figure 5.5: Effect of the field inhomogeneities correction on the fMRI data. Comparison of the mean images(A) and tSNR maps (B) obtained on three volunteers acquired with 3D-SPARKLING at 7T with 1mm isotropic spatial resolution and $TR_{vol} = 2.4s$. In collaboration with Zaineb Amor [Amo+23].

Du passé faisons table en marbre. — Karadoc, Kaamelott

Chapter 6

Benchmarking local low-rank denoising methods for fMRI

Outline

Cor	ntext and Summary	79
6.1	Introduction	80
6.2	Denoising methods using patch based PCA	81
	6.2.1 From global to local low rank assumption.	81
	6.2.2 Local low rank formalism	81
	6.2.3 Comparison of local low-rank methods	82
6.3	Material and Methods	83
	6.3.1 Acquisition	83
	6.3.2 Preprocessing	84
	6.3.3 Statistical Analysis	84
6.4	Results	84
	6.4.1 Single subject analysis	84
	6.4.2 "Group-level" analysis	86
	6.4.3 Noise map estimation	87
6.5	Conclusion	87

The content of this chapter has been presented in an international peer-reviewed conference with proceedings:

P.-A. Comby, Z. Amor, A. Vignaud, and P. Ciuciu. "Denoising of FMRI Volumes Using Local Low Rank Methods". In: ISBI 2023 - International Symposium on Biomedical Imaging. Apr. 18, 2023

Context and Summary

As we presented in Chapters 1 and 2, the acquisition of fMRI data is subject to various types of noise, and a trade-off between spatio-temporal resolution, brain coverage and SNR. This hinders the access to high-quality reference data that could be used for testing the performances of fMRI image reconstruction methods (introduced in Chapters 3 and 4). The initial intent of this work was thus to create such reference data by enhancing (denoising) existing experimental setup.

To this end, we got interested in the Local-Low-Rank denoising methods, among which NORDIC [Viz+21] generated a lot of interest in the community. Yet, we quickly realized that (i) There might be more mathematically optimal methods (ii) The relative position of the denoising step in the preprocessing chain (i.e. before or after the motion correction step and/or the use of magnitude data) may have an impact on the statistical analysis.

To this end, we performed the following benchmark. We conclude that Optimal-Threshold used with complex-valued data is a better than NORDIC, and that applying the motion-compensation step before is a safer choice.

Nonetheless, we came to realize that even with denoising, we would not be able to fully control our reference data to clearly assess the performance of reconstruction methods. Therefore, we put our next effort in the creation of a simulation framework, named SNAKE, presented in Chapter 7.

An interesting further development, would be the use of SNAKE itself to reevaluate our benchmark in an even more controlled environment.

6.1 Introduction

T^{HE} current state-of-the-art of accelerated imaging techniques such as Parallel Imaging[Smi+04] and Compressed Sensing (CS) [Amo+22a] as well as ultra-high magnetic field system (7 Tesla,7 T) [BP11], has enable the collect of fMRI data at submillimetric resolution with reasonable volumic repetition time (TR $\simeq 1$ s to 2 s).

However with such settings, the fMRI data is heavily spoiled by surrounding noise sources: Firstly the thermal noise emanating from the acquisition process contaminates the fMRI signal as a complex independent Gaussian source with spatially varying variance [TPW11]. Secondly the physiological noise, i.e. patient-induced signal variations that are acquired along with the signal of interest but do not reflect significant brain activity (breathing, heart beat, head motion, etc.) [Liu16] (3D-EPI sequences are more sensitive to physiological noise than the classical 2D EPI scheme).

Overall, these noise sources limit the reliability of fMRI data and constrain neuroscientists to collect multiple runs in each participant to boost the statistical sensitivity.

To increase the statistical significance of an fMRI experiment, the sequence of fMRI images is usually pre-processed using state-of-the-art toolboxes such as FSL [Smi+04] and SPM [Fri07]. Recently fMRIprep [Est+19], leveraging the Python based Nipype package [Gor+11] has been proposed as a unified and standardized workflow.

Concomitantly denoising methods have been introduced for fMRI (or other dynamical imaging such as DWI) [Man+13; Ver+16; NNS22; Viz+21; Moe+21; GD14; GD17; JS14; SN10] and all assume Gaussian distributed noise. Yet, their integration in the processing pipeline has hardly been considered, at most the denoising step is performed prior to any processing, or simply tested on synthetic data. Moreover, only the recent contribution of NORDIC [Viz+21], which rooted the interest for this work, considered complex-valued data, which preserve the Gaussian distributed noise hypotheses, instead of using magnitude only images, where the noise distribution is Rician [GP95].

This chapter aims at benchmarking denoising methods for fMRI, by quantifying the gain they provided in the downstream statistical analysis. In particular we evaluated 5 methods based on the local low-rank property of fMRI images (presented in § 6.2) and tested how they can be optimally integrated with data preprocessing. To this end, we acquired 3D-EPI scans in six healthy volunteers undergoing a retinotopic mapping experiment (*cf.* § 6.3), and applied a denoising step either on complex-valued or magnitude-only fMRI images. In § 6.4 the results of standard statistical analysis are summarized over the six individuals. Conclusions are drawn in § 6.5.

6.2 Denoising methods using patch based PCA

6.2.1 From global to local low rank assumption.

Let \boldsymbol{Y} be the *complex-valued* fMRI scans obtained after reconstruction and combination of multi-coil data: $\boldsymbol{Y} = [\boldsymbol{y}_1...\boldsymbol{y}_{N_t}] \in \mathbb{C}^{N_x N_y N_z \times N_t}$, where $N_x N_y N_z$ is the number of voxels and N_t the number of scans. Along with the fMRI data, we can also retrieve a noise variance map $\boldsymbol{\Sigma} \in \mathbb{R}^{N_x N_y N_z}_+$.

FMRI scans are inherently low rank as we observe small variations of the BOLD signal over a constant background. Hence, temporal frames are highly correlated. Furthermore, the periodic behavior of the physiological activities and the presence of repeated stimuli strengthen this effect. In such settings, the data at hand can be considered as a noisy observation of a low rank matrix $\mathbf{Y}_{ij} = \mathbf{X}_{ij} + \mathbf{N}_{ij}$ where the noise components are independent and Gaussian $\mathbf{N}_{ij} \sim \mathcal{N}(0, \mathbf{\Sigma}^2)$. The low rank matrix is the solution of the problem [CST13]:

$$\hat{\boldsymbol{X}} = \underset{\boldsymbol{X}}{\operatorname{arg\,min}} \frac{1}{2} \|\boldsymbol{Y} - \boldsymbol{X}\|_{F}^{2} + \lambda \|\boldsymbol{X}\|_{*}$$
(6.1)

where $\|\mathbf{X}\|_* = \sum_{k=0}^n \sigma_k(\mathbf{X})$ denotes the nuclear norm of \mathbf{X} .

The analytical solution of (6.1) consists in applying a soft thresholding operator to the singular values of \boldsymbol{Y} . However, this global low rank modeling suffers in practice from the high dimensionality of the problem:

- (i) In the case of 3D fMRI acquisitions, we typically have $N_x N_y N_z \sim 10^5 N_t$ and the limited number of degrees of freedom on the SVD will reduce the effectiveness of a rank constraint.
- (ii) The spatial noise level is heterogeneous in the context accelerated imaging due to multi-coil interactions [Bre+09].

To alleviate those challenges, problem (6.1) can be solved locally on small 3D+time patches extracted from the whole 4D fMRI sequence. The patch size and their overlap become hyperparameters for the problem.

6.2.2 Local low rank formalism

The patch extraction operator $\mathcal{P}_{\boldsymbol{u}}(Y)$ basically extracts a $K = k_x \times k_y \times k_z$ dimensional patch centered in voxel \boldsymbol{u} across all fMRI scans in \boldsymbol{Y} , yielding a $K \times N_t$ so-called Casorati matrix (see fig. 6.1). $\boldsymbol{Y}_{[\boldsymbol{u}]} = \mathcal{P}_{\boldsymbol{u}}(\boldsymbol{Y})$ is a low rank matrix containing background (*e.g.a* T_2^* -w version of brain anatomy) information polluted by a zero-mean Gaussian noise of local variance $\sigma_{[\boldsymbol{u}]}^2$: $\boldsymbol{Y}_{[\boldsymbol{u}]} = \boldsymbol{X}_{[\boldsymbol{u}]} + \boldsymbol{N}(\sigma_{[\boldsymbol{u}]}^2)$.

Retrieving a low rank approximation of the patch typically consists in applying a singular value decomposition (SVD) of $Y_{[u]} = USV^T$, and then thresholding it to retain

Chapter 6 Benchmarking local low-rank denoising methods for fMRI

the largest $n_{[u]}$ singular values related to the signal of interest and yield the low rank approximation $\hat{X}_{[u]} = US^*V^T$, cf. fig. 6.1. Once the noise has been discarded, the patches are recombined with a weighting mechanism in case of overlapping patches as originally proposed in [Man+13]. Each voxel at position *i* in the final denoised fMRI sequence, is computed from the *P* patches containing it as follows:

$$\hat{\boldsymbol{X}}(i) = \frac{\sum_{j=1}^{P} w_j \hat{\boldsymbol{X}}_{[\boldsymbol{u}_j]}(i)}{\sum_{j=1}^{P} w_j}, \quad w_j = \frac{1}{1 + n_{[\boldsymbol{u}_j]}}$$
(6.2)



Figure 6.1: General procedure for LLR denoising. Without loss of generality a 2D case is presented. The sequential data is processed into patches, that are (1.) extracted, (2.) SVD decomposed and thresholded, and then recombined (3.) using a weighted average in case of overlap.

6.2.3 Comparison of local low-rank methods

Several local low rank (LLR) denoising methods [Man+13; Ver+16; NNS22; Viz+21; Moe+21; GD14; GD17; JS14; SN10] have been proposed in the literature. We selected four out of them (cf. table 6.2) given their impact in the recent literature and their diversity to address the thresholding problem. When available, we directly used the original authors' implementation. Each method proposed a thresholding function η over the singular values $\lambda_1 \dots \lambda_{N_t}$ of the patch $\mathbf{Y}_{[u]}$. For instance, the optimal threshold (OT) with respect to the Frobenius norm [GD17] is defined as follows:

$$\eta(\lambda) = \frac{N_t \sigma^2}{\lambda} \sqrt{\left(\frac{\lambda^2}{N_t \sigma^2} - \beta - 1\right)^2 - 4\beta} \,\mathbf{1}_{\frac{\lambda}{\sqrt{N_t \sigma}} \ge 1 + \sqrt{\beta}} \tag{6.3}$$

with $\beta = N_t/K$, if a noise map Σ is available, then $\hat{\sigma}$ is the average of $\Sigma_{[u]}$ (referred as Hybrid-OT hereafter), else we resort to the robust estimator $\hat{\sigma} = \text{med}\lambda/\sqrt{N_t\mu_{\beta}}$, and μ_{β} is the median of the Marcenko-Pastur's law [MP67], which describes the asymptotic distribution of the singular values of a normally distributed random matrix of aspect ratio β . This distribution is thus at the heart of the selected methods. Noticeably the adaptive thresholding technique proposed in [JS14] was not considered in the benchmark as its computational time was two order of magnitude larger than its LLR competitors. We also discarded Non-Local-Mean (NLM) based methods [Ber+14; Bhu+16], as they do not fit in the LLR framework. In table 6.2, Hybrid-PCA and Hybrid-OT are variations of MP-PCA and OT approaches, respectively, in which a noise map (i.e. a covariance matrix Σ) estimate was provided.

Table 6.2: LLR methods under study. Defining the threshold for each patch in Hybrid-PCA requires an external noise variance map, not required by MP-PCA. NORDIC uses a global threshold, and normalizes the noise variance of each patch to apply it. Optimal-Threshold is described in more detail in text, cf eq. (6.3). $\theta_{[u]}$ denote a patch-local threshold and θ_{global} a threshold value common to all patches.

Name	Thresholding $\eta(\lambda)$	Extra Data used
Nordic [Viz+21]	$\max(0, \lambda - \theta_{global})$	None
MP-PCA [Ver+16]	$\max(0, \lambda - \theta_{[\mathbf{u}]})$	None
Hybrid-PCA [NNS22]	$\max(0, \lambda - \theta_{[\mathbf{u}]})$	Σ
OT [GD17; Zhu+22]	eq. (6.3)	None
Hybrid-OT [GD17]	eq. (6.3)	Σ

For ease of comparison, we settle a patch size of $11 \times 11 \times 11$ voxels with an overlap of 5 voxels in each direction. Adding overlap reduces boundary effects, at the cost of larger processing time. In order to mitigate this effect, we only computed patches that shared at least 10% of their voxels with the brain mask, halving roughly the computational time.

6.3 Material and Methods

6.3.1 Acquisition

Six healthy volunteers were scanned on a Siemens Magnetom 7 T (Siemens-Healthineers, Erlangen, Germany) and a 1Tx-32Rx head coil (Nova Medical, Wilmington, CO, USA) using a 3D EPI sequence (1 mm - iso TE=20 ms, volumetric TR=2.4 s, 120 reps). Taskbased fMRI data was collected along two consecutive runs during a retinotopic mapping paradigm, implemented in ¹, with a rotating wedge (clockwise and anti-clockwise) with a period of 32 s [22]. This paradigm promotes well localized BOLD signal in the visual areas, however the spatial and temporal resolution considered were challenging in regard to the temporal SNR. The LLR denoising methods should thus bring significant benefit to the downstream task of fMRI data analysis in order to detect evoked brain activity and reconstruct the retinotopic phase maps.

¹https://github.com/hbp-brain-charting/public_protocols

6.3.2 Preprocessing

To quantify the potential benefit of each denoising method we tested 4 distinct preprocessing pipelines that consist in applying the denoising step either on the complex-valued (CD) or magnitude-only (MD) fMRI images and then in interchanging its position in the workflow (i.e. before or after image realignment (R) for motion correction). In case of complex denoising after realignment (i.e. R+CD scenario) the motion correction estimated from the magnitude images was applied to the real and imaginary parts before denoising. in the opposite scenario (CD+R), we computed the magnitude of fMRI images prior to performing motion correction.

Additionally, a denoise-only scenario was tested, but it underperformed compared to the standard workflow that embeds realignment. In every case the fMRI images are corrected for B_0 field inhomogeneities afterwards. To favor reproducible research, we provide open sourced implementations of the LLR denoising methods² and of the preprocessing workflow ³.

6.3.3 Statistical Analysis

After testing the different configurations of the preprocessing pipeline, the fMRI images data were analyzed using the Nilearn package [Abr+14]. First, a general linear model (GLM) that embodies the two runs associated with the retinotopic experimental paradigm (clockwise and counter-clockwise) was built up. It includes 2 paradigm-related regressors (parametric, continuous and sinusoidal), 6 rigid motion regressors, a drift regressor and the baseline. The global effect of interest was determined using a statistical *F*-test over the two sinusoidal regressors ($H_0: \alpha_{1,i}^2 + \alpha_{2,i}^2 = 0$), providing a z-score in each voxel *i*. To detect evoked brain activity we then thresholded the z-score map at p < 0.05, to correct for multiple comparisons using the false discovery rate (FDR).

6.4 Results

6.4.1 Single subject analysis

In fig. 6.3, we compare the z-score distributions related to the global effect of interest associated with the original and denoised fMRI images for a single subject. LLR desnoising methods (color-coded) reach higher z-scores compared to the noisy baseline (in gray) whatever the type of fMRI images processed (magnitude-only or complex-valued). We selected this subject as that showing the best improvement between the baseline and the denoised versions.

It is worth noting that the type of input fMRI images slightly impacts the tail of the z-score histograms (see fig. 6.3) as the noise statistics becomes Rician for magnitude-only images.

Moreover, the order of preprocessing steps (R: realign, D: denoising) matters when applied to complex-valued fMRI images: In that context, denoising must precede realignment as this combination provides the largest gain in statistical sensitivity. In contrast, for

 $^{^{2}} https://github.com/paquiteau/patch-denoising$

³https://github.com/paquiteau/retino-pypeline

this subject the best order of preprocessing steps applied to magnitude-only fMRI images is less clear. Hence, to obtain more reliable results, we performed the same analysis on the 6 participants. However, prior to analyzing the results at the group level (cf. § 6.4.2), we pay attention to the spatial effect of denoising on the statistical z-score maps for this subject.



Figure 6.3: Non-zero Z-score probability distribution (ν) for a single participant (subject-3). Top vs Bottom: Analysis of magnitude-only vs complex-valued fMRI images. The distribution has been smoothed with a Gaussian kernel for visualization purposes.

In fig. 6.4 we show the denoised z-score maps yielded by the 20 tested pipelines as well as the noisy baseline. Since we only report surviving voxels to thresholding, we bring evidence that NORDIC, MP-PCA and Hybrid-PCA produced more false positives as they retrieve activations in the white matter on top of the expected ones in the visual cortex. Indeed the retinotipic paradigm is known to elicit evoked activity only in the occipital cortex.

Hence, OT-based approaches are more specific as their activations remain located in the gray matter. At the same time, the OT approach yields the highest z-scores with the largest spatial extent in the occipital region (cf fig. 6.5), notably when the denoising is performed prior to realignment (scenario OT/MD+R). In this subject, we noticed that the behavior of Hybrid-PCA and MP-PCA is similar, and those approaches underperform compared to NORDIC and OT-based methods.





Figure 6.4: Z-score activation maps (axial slices) for the same participant as in fig. 6.3. Only non-zero values are displayed.

In fig. 6.5 we actually zoomed in the occipital cortex and show first that all LLR denoising methods increase the spatial extent and significance of z-scores, hence confirming the whole brain results depicted in fig. 6.3. Further, we demonstrate that the OT approach provides the best sensitivity/specificity trade-off even though we cannot access to the ground truth as we are not in a simulated framework.

6.4.2 "Group-level" analysis

In table 6.6 we report the average gain across subjects for all tested denoising pipelines as a multiplicative factor MF_s obtained by dividing for each subject s the number of activating voxels associated with each LLR method over the count of activating voxels for the baseline (no denoising). Then we averaged these factors across subjects to get $\overline{MF} = 1/6 \sum_{s=1}^{6} MF_s$. We computed this multiplicative factor first over the whole brain but only for the voxels surviving to FDR-corrected thresholding and then within a region of interest (ROI) located in the occipital cortex.

When comparing columns in table 6.6, we can see that in almost every case, performing the denoising step prior to realignment is the most beneficial whether this treatment is applied to magnitude-only (MD+R) or to complex-valued (CD+R) fMRI images. This result is particularly significant when focusing on the occipital ROI.

Even though the realignment step aims at producing lower-rank fMRI images notably by removing part of the physiological noise, this step also resamples the images, thus potentially affecting the noise behavior and making it no longer additive. This could explain why the denoising step is less efficient when performed after motion correction, i.e. in R+MD vs MD+R or in R+CD vs CD+R.

In contrast to the claims done in [Moe+21; Viz+21], our benchmark summarized in table 6.6 shows that NORDIC on average possesses the smallest gain in activation (smallest \overline{MF} value) compared to the baseline and its competitors. This is notably true when the denoising step is applied to complex-valued fMRI images. However, when considering magnitude-only images as input parameters to denoising, Hybrid-PCA on average is the best performer, whether we look at the whole brain or ROI-restricted statistical analysis. This demonstrates that the behavior of this approach reported for subject 3 in § 6.4.1 is not fully representative of the 6 participants.

6.4.3 Noise map estimation

MP-PCA and Hybrid-PCA give access to an a posteriori estimate of the noise variance map (not shown). Firstly, this extra information can be used by NORDIC [Moe+21], which takes as optional input parameter a *g*-factor map. The latter is usually cumbersome to compute as it necessitates to collect additional noise-only data. However, once computed, this map can be used to perform a spatial normalization.

Secondly, when considering MP-PCA or OT approaches that do not rely on this extra knowledge by default, this noise map could be further used in the statistical model for fMRI data analysis to solve a weighted ordinary least squares (OLS) when fitting the GLM parameters instead of a single OLS.

6.5 Conclusion

In this study, we tested multiple LLR denoising methods as an additional step of high resolution fMRI images preprocessing. We demonstrated the added value of injecting this step as early as possible in the workflow in terms of gain in statistical sensitivity. Among the 5 tested methods, we found that the Optimal Threshold approach provides the best sensitivity-specificity trade-off as it does not retrieve false positives in the white matter while boosting the statistical significance in the visual cortex. Further we also realized there is not a *one size fits all solutions*: While the OT method yields the best results when performing the denoising step on complex-valued fMRI images (which are not always available), Hybrid-PCA is the most promising when denoising magnitude-only images. This preliminary study on task-based fMRI calls for a broader validation, both on a larger cohort and in other use cases, notably in resting-state fMRI. Importantly, the use of extra data such as a noise map in Hybrid-PCA and Hybrid-OT is not mandatory to already obtain a significant improvement: Sometimes it helps (Hybrid-PCA vs MP-PCA), sometimes it does not (Hybrid-OT vs OT). Future work will also involve analyzing the impact of hyperparameters of denoising methods (e.g. the patch size and their overlap) on the downstream task.



Figure 6.5: Zoom in the occipital cortex of z-score activation maps (axial slices) for the same participant as in fig. 6.3, thresholded at p < 0.05 using FDR correction for multiple comparisons.

Table 6.6: Average gain (multiplicative factor) \overline{MF} of activated voxels using denoising methods compared to baseline. Over 6 subjects and test cases, the highest gain is achieved by OT (**x8.03**) in the R+CD configuration. Whole brain vs ROI-based \overline{MF} are reported as top/bottom numbers in each cell of the table.

Denoiser	NORDIC	MP-PCA	Hybrid-PCA	OT	Hybrid-OT
R+MD	$ imes 3.52 \ imes 3.64$	${ imes 6.02} { imes 4.92}$	$\substack{\times 6.09\\\times 4.93}$	×3.73 ×4.33	${ imes}0.91 { imes}1.00$
R+CD	${ imes}0.57 \\ { imes}0.55$	${ imes 2.98} \\ { imes 2.70}$	$\substack{\times 3.27\\\times 2.91}$	$\begin{array}{c} \times 8.03 \\ \times 5.04 \end{array}$	$ imes 5.29 \ imes 4.01$
MD+R	$ imes 3.36 \ imes 3.48$	$ \times 6.32 $	$\begin{array}{c} \times 7.55 \\ \times 6.10 \end{array}$	×3.10 ×3.19	$^{\times 1.22}_{\times 4.90}$
CD+R	${ imes 2.57} { imes 2.59}$	$ imes 5.53 \\ imes 4.39$	$\substack{\times 4.97\\\times 4.27}$	$\begin{array}{c} \times \textbf{7.91} \\ \times \textbf{6.26} \end{array}$	×5.45 ×5.00

Lourd est le parpaing de la réalité sur la tartelette aux fraises de nos illusions. — Boulet

Chapter 7

No Ground-Truth ? Let's build some

Outline

Con	ntext and Summary	92
7.1	Motivation	93
7.2	Existing software and their limitations	94
7.3	Extended Fourier model for fMRI acquisition	96
	7.3.1 Single shot acquisition	96
	7.3.2 BOLD as a TE-sensitive change between shots	98
	7.3.3 Noise and SNR calibration	98
	7.3.4 Summary of general hypotheses on the acquisition model 1	00
7.4	Main characteristics of SNAKE implementation	00
	7.4.1 Modular Approach	00
	7.4.2 Performance, reproducibility and scalability made easier for	
	neuroscientists	02
7.5	Numerical experiments	03
	7.5.1 Acquisition scenarios	03
	7.5.2 Reconstruction Strategies for Scenario 2 & 3	06
	7.5.3 Evaluation methods $\ldots \ldots 1$	07
7.6	Results	07
	7.6.1 Scenario S1	07
	7.6.2 Scenario S2	10
	7.6.3 Scenario S3	13
7.7	Discussion	13
	7.7.1 Effect of T_2^* relaxation $\ldots \ldots \ldots$	13
	7.7.2 Setting the noise level and tSNR	14
	7.7.3 Insights from the different scenarios	14
	7.7.4 Limits to the study $\ldots \ldots \ldots$	15
	7.7.5 Extending the simulator	16
	7.7.6 Exploring the effect of tuning acquisition and reconstruction	
	together $\ldots \ldots 1$	17
	7.7.7 More than a fMRI simulation tool	18
7.8	Conclusion	18

The content of this chapter has been submitted for publication:

P.-A. Comby, A. Vignaud, and P. Ciuciu. *SNAKE-fMRI: A Modular fMRI Data Simulator from the Space-Time Domain to k-Space and Back*. Version 2. Apr. 12, 2024. URL: http://arxiv.org/abs/2404.08282 (visited on 09/01/2024). In revision to Imaging Neuroscience

It also has been presented as an abstract at ISMRM 2024:

P.-A. Comby, A. Vignaud, and P. Ciuciu. "SNAKE-fMRI: A Modular fMRI Simulator from the Space-Time Domain to k-Space Data and Back". In: *ISMRM Annual Meeting, (in Press).* Singapore, 2024. Awarded the 2nd best abstract award from the reproducibility study group.

Context and Summary

In Chapters 3 and 4 with highlighted that the access to proper reference data was of paramount importance for validating fMRI image reconstruction methods. However, getting such data from *in vivo* experiments is in practice unfeasible, even if we conducted perfect processing (akin to the denoising in Chapter 6), as we cannot provide absolute guarantee for the true location of the activation.

Instead, we opted to create an *in silico* setup by modeling the whole fMRI models that we introduced in Chapter 1, as well as the technological improvements presented in Chapter 2.

We named this simulation framework SNAKE: Simulator from Neuro-Activation to K-space Evaluation. It is a open-source, modular and extensible Python-based simulator capable of efficiently generating high-resolution fMRI raw k-space data from evoked brain responses using various 3D sampling strategies of k-space data and multiple coils¹. Additionally, we also model the T_2^* relaxation effect. as it plays a significant role in real fMRI acquisition, impacting the statistical power in challenging acquisition scenarios.

Using this *in-silico* setup, we can provide a realistic and reproducible ground truth for fMRI reconstruction methods in 3D accelerated acquisition settings and explore the influence of critical parameters. This includes the acceleration factor and signal-to-noise ratio (SNR), on downstream tasks of image reconstruction and statistical analysis of evoked brain activity.

We present three scenarios of increasing complexity to showcase the flexibility, versatility, and fidelity of SNAKE: From a temporally fixed full 3D Cartesian to various 3D non-Cartesian sampling patterns, we can compare — with reproducibility guarantees — how experimental paradigms, acquisition strategies and reconstruction methods contribute and interact together, affecting the downstream statistical analysis.

These scenarios are reconstructed using a Compressed Sensing frame-wise method $(\S 3.1.2)$ with a warm start setup. Our results show that, when combined with a time-varying acquisition strategy, the frame-wise + warm start combination provides a good image quality and good statistical performances compared to the classical approaches.

SNAKE paves the way for the development of new data acquisition and image reconstruction strategies for fMRI, and notably in the context of deep-learning.

¹using the extended Fourier Model provided by MRI-NUFFT, see Chapter 5

7.1 Motivation

Conducting task-based fMRI experiments is an expensive and time-consuming endeavor, often requiring access to advanced imaging facilities and substantial expertise in protocol design, data collection, and analysis. In addition, the reproducibility of fMRI experiments and findings has been shown to be a critical issue in cognitive and clinical neuroscience [Lio+06; Gri+16; Nee19; Bos+20].

However, a major roadblock to replicate fMRI studies is the sample size, as most studies involved small cohorts and are less replicable mainly due to multiple sources of variability (inter-subjects, MR systems, etc. Nakuci et al. [Nak+23] and Marek et al. [Mar+22]). Moreover, accommodation and habituation effects limit the repeatability of the intra-subject test [BM10], which severely hinders the comparison of acquisition protocols.

This variability prevents the development of new methodologies to answer the growing needs in the fMRI literature, particularly the quest for higher temporal [VP21] and spatial [Pin+20] resolution. The development of complex methods on data acquisition and image reconstruction in (f)MRI faces challenges in providing a fair comparison due to limited existing ground truth data collected on the same individuals under different imaging setups. Recent advancements in computing power have led to the emergence of (f)MRI simulators, which aim to address this limitation (see Welvaert and Rosseel [WR14] for a review). However, most current simulators are restricted to producing magnitude-only (f)MRI images and lack raw k-space data, severely hindering their usability for validation of acquisition and reconstruction methods. On the other hands, Anatomical MRI simulators typically relies on the Bloch's model [Cas+23], but are too computationally demanding and lack parametrization for the fMRI experiment.

In the search for an absolute ground truth, other solutions than simulation have been proposed, such as active phantoms [Che+06; Kum+21], fMRI monitoring with other devices [CR17], or removing the acquisition artifacts [Amo+24a; Due+15; Bol+17]. More recently, the development of generative Artificial Intelligence (AI) has opened new opportunities for the synthesis of realistic MRI data [Pin+23; Gop+24]. However, it requires expensive computation that would not scale for generating high-resolution 3D+time fMRI data, and the lack of control and explicability over the yielded sequence of fMRI volumes.

In general, none of the proposed solutions has succeeded in providing a (i) a full control on both the temporally resolved input hemodynamic signal and the output BOLD fMRI time series, (ii) an easy-to-use and reproducible framework through open-source software and (iii) a low computational or operational cost.

This chapter aims to fill this gap. More precisely, our contribution can be summarized as follows:

We propose a realistic fMRI simulator, named SNAKE based on an *extended* Fourier model of MRI data acquisition that can create all the required k-space data for evaluating the fMRI processing tool chain: From the definition of an experimental paradigm and the localization of brain activation in a realistic phantom, up to the generation of 3D+time k-space data. This bottom-up approach gives fMRI scientists a level playing field to explore and compare various properties of acquisition and reconstruction strategies, such as the choice of the sampling pattern (Cartesian vs. non-Cartesian readouts), the acceleration factor, the signal-to-noise ratio (SNR), etc. SNAKE could as well be used in the near

future to train deep learning models for fMRI image reconstruction, by providing fully parameterized high-quality reference data.

To compare the reconstruction methods in more challenging settings, the ground truth data can be degraded in various ways before and during the acquisition process (motion, static field inhomogeneities, thermal and physiological noise, etc.).

The simulator primarily produces a sequence of k-space volumes sampled with realistic fMRI sequences. Furthermore, for completeness, we offer the possibility to produce a sequence of fMRI volumes in the image space by plugging in simple but efficient *volume-wise* reconstruction algorithms. SNAKE has been designed to be as light and as fast as possible to give its users an upper limit to the statistical results that an fMRI experiment can achieve given the acquisition and reconstruction strategies chosen.

In the remainder of this chapter, we first propose a review of the available fMRI simulation tools, highlighting their strengths and weaknesses (§ 7.2). Then, we present the underlying model (§ 7.3) and the design of SNAKE (§ 7.4). Finally, we present three typical simulation scenarios (§ 7.5), with basic image reconstruction and statistical analysis (§ 7.6), to illustrate the various possibilities of SNAKE, and discuss further the strength and limitations of § 7.7.

Along with this chapter, we also provide a detailed documentation and a set of tutorials to help users get started with SNAKE in the accompanying documentation available at https://mind-inria.github.io/snake-fmri.

7.2 Existing software and their limitations

The intricacies of fMRI data create significant obstacles in developing a unified framework for generating comparable datasets. A detailed examination of the existing literature on fMRI simulation, as highlighted by [WR14], has shown that the lack of a standardized approach to synthesizing fMRI data severely hinders reproducibility in fMRI research. This underscores the need for improved transparency in the reporting of experimental design and a more nuanced understanding of the processes involved in the acquisition of fMRI data. In addition, there exist anatomical MRI simulators based on the Bloch equations that generate contrast-weighted MR images, but they usually remain confined to anatomical imaging with no straightforward extension to fMRI. A detailed comparison of the current available MRI simulators is provided in [Cas+23].

In what follows, we provide a broad overview of available MRI and fMRI simulators in Table 7.1 and focus hereafter on the main alternative to SNAKE, in increasing order of similarity.

Among existing fMRI data simulation tools, several offer distinct advantages and limitations. The Virtual Brain [Sch+22; San+13] is an open-source multi-modal brain simulator, focusing on the simulation of brain activity using the structural connectome and interaction between functional networks. It is capable of yielding voxel-wise BOLD fMRI time series. However, it remains complex to master, and is not a dedicated tool for analyzing the acquisition/image reconstruction chain in fMRI.

fMRIsim proposed by Ellis et al. [Ell+20], is a Python package that enables standardized and realistic fMRI time series simulation in the image domain. It was inspired by a parent R package neuRosim [Wel+11]. It allows for the evaluation of complex experimental designs and optimization of the statistical power. However, it focuses on single-subject

	Simulator Name	Source	API	Sim. Do- main	Required External Data	Interface	Reconstr.	Ecosysten
	The Virtual Brain [San+13]	GPL- 3.0	Ç	Image		GUI/scril	N/A	N/A
	Jemris [Sto+10]	GPL- 2.0	4 ©	Bloch		GUI	ISMRMD raw data	N/A
or	ODIN [Joc+06]	GPL- 2.0	3	Bloch	Tissue Maps, Sequence	c++/GU.	\mathbf{FFT}	OdinReco
Simulat	MRILab [Liu+17]	BSD-2	1	Bloch	Preset Macros	GUI	FFT Non- Cartesian	GUI Gadgetroi
MRI	Bloch-Solver [KK17]	Proprieta	e	Bloch	Tissue Maps,	script	\mathbf{FFT}	N/A
	Fabian [Laj+22]	BSD-3	-	Bloch (EPG)	Tissue Maps, Sequence	script	\mathbf{FFT}	N/A
	KomaMRI [Cas+23]	MIT	••	Bloch	Pulseq	GUI	\mathbf{FFT}	Pulseq
	POSSUM [Dro+06]	FSL	10	Bloch	Tissue Maps Sequence, Events	CLI	FFT	FSL
ulato	Neurolib [CJO23]	MIT	¢	Image	Connectivity Matrices	script	N/A	Jupyter
MRI Sim	SimTB [Erh+12]	Open Source	-	Image	Spatial Maps, Events	GUI	N/A	MATLAB
ţ	NeuRoSim [Wel+11]	GPL- 2.0	R	Image		script	N/A	N/A
	fmriSim [Ell+20]	Apache- 2.0	¢	Image		script	N/A	Brainiak
	SNAKE-fMRI	MIT	ę	Kspace+	Configuration files	script/CL	Any	Any

Table 7.1: Summary of characteristics of different fMRI data simulators.

simulations and requires manual parameter setting or estimation from real data. Moreover, it primarily deals with magnitude data and uses additive noise settings, which might restrict its applicability to specific simulations or use cases (e.g. limited resolution).

SimTB, introduced by Erhardt et al. [Erh+12], is a MATLAB toolbox specialized in simulating fMRI time series using a separable spatio-temporal model. It offers extensive customization options, including spatial sources, experimental paradigms, tissue-specific properties, noise, and head movement. SimTB is equipped with both a graphical user

interface and scripting capabilities. However, no k-space data is available to assess the performances of various image reconstruction algorithms.

The POSSUM simulator, as outlined by Drobnjak et al. [Dro+06], offers a comprehensive approach to the impact of specific artifacts encountered in the acquisition of fMRI data. POSSUM accurately simulates these artifacts using Bloch equations and a geometric definition of the brain. However, its computational cost is high, making simulations of full brains at high resolution prohibitive. Additionally, it currently only offers the possibility to yield Cartesian data in k-space. FMRI analysis is deferred to the FSL library, and no comparison with the ground truth is provided as an outcome of the toolbox.

7.3 Extended Fourier model for fMRI acquisition

To bridge the gap between image and k-space simulation, a more realistic forward model is needed. Using the Bloch equation model [Blo46] would be prohibitively expensive for whole-brain simulation at high spatial and temporal resolution; instead, we extend the classical model based on the Fourier transform.

7.3.1 Single shot acquisition

Considering the multicoil imaging setup, T_2^* relaxation and off-resonance effects due respectively to signal decay and B_0 inhomogeneities over the field of view (FOV), the signal acquired by the MR system for the s-th shot in the ℓ -th coil is defined as follows:

$$y_{\ell,s}(t) = \int_{FOV} m(\boldsymbol{r}, t) \, \mathcal{S}_{\ell}(\boldsymbol{r}) \, e^{-2\imath \pi \Delta f_{r}(\boldsymbol{r})t} \, e^{-2\imath \pi \boldsymbol{k}_{s}(t) \cdot \boldsymbol{r}} d\boldsymbol{r}$$
(7.1)

where \mathbf{k}_s is the k-space trajectory for the s-shot with $k_s(t)$ the k-space location at time t, $\Delta f_r(\mathbf{r})$ the static B_0 inhomogeneity map, $\mathcal{S}_{\ell}(\mathbf{r})$ the sensitivity map for coil $\ell \in \{1, \ldots, L\}$, $m(\mathbf{r}, t)$ the base magnetization of the image. Note that Eq. (7.1) holds for both 2D and 3D imaging. However, the main focus of this chapter will be 3D acquisitions, which can proceed to segment the readout either across planes for stacked strategies (e.g. 3D EPI, stack of spirals) [Pos+10] or across shots for full 3D non-Cartesian strategies (e.g. SPARKLING or TURBINE, Graedel, Miller, and Chiew [GMC22]). In the former case, we will use $k_s(t) = [k_s^x(t), k_s^y(t)]$ at a given elevation z-th embedded in the definition of s, whereas in the latter case we will consider $k_s(t) = [k_s^x(t), k_s^y(t), k_s^z(t)]$.

Assuming steady-state and perfect spoiling regimes, the signal obtains from a Gradient Recall Echo (GRE) sequence parameterized by (TR, TE, α) , with $T_1 \gg TR$, is:

$$m(\mathbf{r},t) = \rho(\mathbf{r}) \sin \alpha \, \frac{1 - e^{-TR/T_1(\mathbf{r})}}{1 - \cos \alpha \, e^{-TR/T_1(\mathbf{r})}} \, e^{-t/T_2^*(\mathbf{r})} = \mu(\mathbf{r},t_{ref}) e^{(t_{ref}-t)/T_2^*(\mathbf{r})} \tag{7.2}$$

where $\rho(\mathbf{r}), T_1(\mathbf{r}), T_2^*(\mathbf{r})$ are quantitative maps of the proton density, longitudinal (T_1) and transverse (T_2^*) relaxation parameters, and have fixed values during the shot. Different contrasts can be generated, such as the one shown in Figure B.1. $\mu(\mathbf{r}, t_{ref})$ describes the state of the object (i.e. its contrast) obtained after the excitation pulse at a point of reference with respect to relaxation t_{ref} , in the case of the in-out sampling pattern, we typically have $t_{ref} = TE$, the echo time centered on the T_{obs} long read-out window with $t \in \left[-\frac{T_{obs}}{2}, \frac{T_{obs}}{2}\right]$. We note
To simplify the computation, the quantitative maps are separated in N_{tis} tissue types each with a set of constant MR parameters T_1 , T_2^* and ρ :

$$m(\mathbf{r},t) = \sum_{i=1}^{N_{tis}} m_i(\mathbf{r},t) = \sum_{i=1}^{N_{tis}} w_i(\mathbf{r}) \,\mu_i(t_{ref}) \,e^{(t_{ref} - t)/T^*_{2,i}}$$
(7.3)

where w_i is the proportion of tissue *i* at voxel location \mathbf{r} , $\mu_i(t_{ref})$ the contrast of the *i*-th tissue at time t_{ref} (which depends the tissue's $T_{1,i}$) and $T_{2,i}^*$ the T_2^* relaxation parameter of the *i*-th tissue type.

Combining eqs. (7.1) to (7.3), we obtain the complete signal equation for a single coil ℓ and a single shot s:

$$y_{\ell,s}(t) = \int_{FOV} \sum_{i=1}^{N_{tis}} w_i(\boldsymbol{r}) \, \mu_i \, e^{-t/T_{2,i}^*} \, \mathcal{S}_{\ell}(\boldsymbol{r}) \, e^{-2i\pi\Delta f_r(\boldsymbol{r})t} e^{-2i\pi\boldsymbol{k}_s(t)\cdot\boldsymbol{r}} d\boldsymbol{r}$$

$$= \sum_{i=1}^{N_{tis}} \mu_i \, e^{-t/T_{2,i}^*} \int_{FOV} w_i(\boldsymbol{r}) \, \mathcal{S}_{\ell}(\boldsymbol{r}) \, e^{-2i\pi\Delta f_r(\boldsymbol{r})t} e^{-2i\pi\boldsymbol{k}_s(t)\cdot\boldsymbol{r}} d\boldsymbol{r}.$$
(7.4)

The static off-resonance $e^{-2i\pi\Delta f_r(r)t}$ term can be approximated by a sum of separable bilinear terms $\sum_{p=1}^{P} c_p(t) b_p(r) \simeq e^{-2i\pi\Delta f_r(r)t}$ [SNF03; Amo+23], where *P* is the number of interpolator used to approximate the decomposition.

Furthermore, considering the set of k-space sampling points $\mathbf{k}_1, \ldots, \mathbf{k}_n, \ldots, \mathbf{k}_N$ that are acquired possibly off the Cartesian grid and at times $t_n = (n-1)\Delta t$ (where Δt is typically the dwell time of the scanner, in the order of 10μ s), and the spatial locations $(\mathbf{r}_m)_{m=1}^M \in \mathbb{N}^{3 \times M}$ in $M = N_x N_y N_z$ voxels to cover the FOV, we obtain:

$$y_{\ell,s}[t_n] = \sum_{i=1}^{N_{tis}} \mu_i \, e^{-t_n/T_{2,i}^*} \sum_{p=1}^{P} c_p[t_n] \, \int_{FOV} w_i[\mathbf{r}] \mathcal{S}_{\ell}(\mathbf{r}) b_p(\mathbf{r}) e^{-2i\pi \mathbf{k}_s[t_n] \cdot \mathbf{r}} d\mathbf{r}$$
$$= \sum_{i=1}^{N_{tis}} \mu_i \, e^{-t_n/T_{2,i}^*} \sum_{p=1}^{P} c_p[t_n] \, \sum_{m=1}^{M} w_i[\mathbf{r}_m] \, \mathcal{S}_{\ell}[\mathbf{r}_m] \, b_p[\mathbf{r}_m] \, e^{-2i\pi \mathbf{k}_s[t_n] \cdot \mathbf{r}_m}$$
$$= \sum_{i=1}^{N_{tis}} \mu_i \, e^{-t_n/T_{2,i}^*} \sum_{p=1}^{P} c_p[t_n] \, \mathcal{F} \left\{ b_p \, \mathcal{S}_{\ell} \, w_i \right\} [\mathbf{k}_s[t_n]] \,.$$
(7.5)

The resulting signal is a doubly weighted sum of Fourier transforms (\mathcal{F}) first by interpolating coefficients in the k space (b_p) and second by tissue-specific contrast and T_2^* decay: $\mu_i e^{-t_n/T_{2,i}^*}$. In general, to generate the k space data $\boldsymbol{y} = (y_1, \ldots, y_L)$ in the multicoil array, for a single shot s, the total number of Fourier transform calls in (7.5) is $N_{tis} \cdot P \cdot L$. The computational cost may be reduced by limiting the number of tissues to a few (e.g. 3 like gray matter, white matter, and cerebrospinal fluid) and / or neglecting T_2^* relaxation (which is admissible if $T_{obs} = N\Delta T \ll T_2^*$) and off-resonance effects. Making all these hypotheses lead to the following *Basic Fourier* model:

$$y_{\ell,s}[t_n] = \mathcal{F}\left\{\mathcal{S}_{\ell}\sum_{i=1}^{N_{tis}} \mu_i w_i\right\} [\mathbf{k}_n] = \mathcal{F}\left\{\mathcal{S}_{\ell} \,\boldsymbol{\mu}\right\} [\mathbf{k}_s[t_n]].$$
(7.6)

Note that (7.6) is the model used as the basis for the reconstruction algorithms. However, neglecting T_2^* relaxation in fMRI acquisition may yield misleading results, in particular for long-readout trajectories (such as EVI) or if two neighboring points in the k-space are sampled at the two extremities of the sampling trajectory.

7.3.2 BOLD as a TE-sensitive change between shots

The previous section introduced the acquisition model for a single shot. In practice, acquisition consists of multiple shots, acquired at every TR_{shot} . In the context of 3D fMRI, these shots are then grouped together to build a single k-space volume (see Figure 7.2). Adding functional MRI capability is done by modifying simulation parameters between shots. In particular, the BOLD effect is modeled as a change of T_2^* following If we consider that the BOLD effect modifies the baseline gray matter T_2^* — in a given region of interest, and in a simplified manner — as $T_{2,BOLD}^* = \frac{1}{R_{2,GM}^* + \Delta R_2^*}$ [Jin+06; UK23], we have the following:

$$\mu_{BOLD} = S_0 \cdot \exp(-TE/T_{2,BOLD}^*) = S_0 \cdot \exp(-TE \cdot (R_{2,GM}^* + \Delta R_2^*))$$

$$\mu_{BOLD} = \mu_{GM} \exp(-TE \cdot \Delta R_2^*) \simeq \mu_{GM} (1 - TE \cdot \Delta R_2^*).$$
(7.7)

where S_0 denotes the net magnetization Hence, for every shot s, we can determine the base intensity, in the image domain:

$$\mu_{BOLD}(t_s) = \left(1 - TE \cdot \Delta R_2^* \tilde{h}(t_s)\right) \cdot \mu_{GM}(t_s), \qquad (7.8)$$

where $\tilde{h}(t)$ is the normalized hemodynamic response such that $\max{\{\tilde{h}(t)\}} = 1$. In the context of task-based fMRI, it is modeled by convolving a sequence of events (event-related paradigms) or blocks (block paradigms) with a reference hemodynamic response function HRF [ciuciu2003unsupervised; Glo99]. In our simulation, we used $\Delta R_2^* = -1$ Hz (following the value used by Jin et al. [Jin+06]), which generates a 2. 5% increase in the BOLD contrast at TE = 25ms.

7.3.3 Noise and SNR calibration

The (f)MRI signal in the k-space is also corrupted by thermal noise sources that arise from the acquisition process in two forms: Brownian motion of spins and random fluctuation in the RF receiver processing chain.

To model those effects, we add a complex multivariate Gaussian noise over the coil component $\ell \in (1, \ldots, L)$ for every shot s. An existing coil covariance matrix $\Sigma \in \mathbb{C}^{L \times L}$ can be supplied to match an existing hardware set-up. Moreover, the global noise variance can be tuned manually, to set the *input SNR* in k-space. Concretely, the noise for each time point in the shot is sampled from $(n_{\ell,s}[t_n])_{1 \leq \ell \leq L} \sim \mathcal{N}(0, \frac{E(\hat{x})}{SNR_i}\Sigma)$ where Σ is the coil covariance matrix. $E(\hat{x})$ is the energy of the *ideal* phantom, acquired at TE, and SNR_i is the input SNR defined by the user in k-space. To add noise to a full shot, we draw Nrealizations of the L-dimensional noise vector $(n_{\ell,s}[t_n])_{1 \leq \ell \leq L}$. Then noisy data is formed $1 \leq n \leq N$

as follows:

$$\forall t_n \in 1, \dots, N, \quad \tilde{y}_{\ell,s}[t_n] = y_{\ell,s}[t_n] + n_{\ell,s}[t_n].$$
(7.9)

It is possible to calibrate the value of SNR_i using experimental data, by computing the energy ratio of k-space shots collected with and without RF excitation.

Another solution is to view the joint system of acquisition and reconstruction together, and consider a known case of image quality output. In this case, the value of SNR_i is tuned to obtain the desired image quality, gathering all missing modeling aspects in this Gaussian noise.

More structured noise sources (e.g. physiological noise such as heartbeat, breathing rate, motion) could be superimposed on the simulated signal. However, the wide variety of options in this field for modeling such noise components is beyond the scope of this article. However, SNAKE already offers a wide flexibility through the handler mechanism (see § 7.4.1). Future work (open to contributions) will go into more detail on the implementation of these noise sources.



Figure 7.2: Acquisition method implemented in SNAKE – The case represented is simplified to a 2D Cartesian case (*e.g.* a projected view of a 3D non-accelerated EPI scheme). Each shot (i.e. a plane in 3D EPI) of the k-space sampling pattern is acquired separately from an on-the-fly simulated volume in the image domain as shown in the blue frame. The shots are numbered here from ① to ⑦. The parallel acquisition is performed in parallel for each tissue type to apply the T_2^* relaxation model (7.5).

Extending the signal model

In § 7.3 we presented the acquisition model of SNAKE. And in § 7.3.2 how the BOLD signal is updated between the acquisition of two shots. Similarly, other updates of the acquired data could happen between shots, such as motion, physiological noise, or other perturbations, that will be modeled using the handler mechanism (as outlined in Figure 7.3). For instance, we present the effect of motion as an example in the

accompanying software documentation². Concretely, for the motion case, we would update the position of all the tissues masks w_i (i.e. with a temporal resolution of $TR_{shot} \simeq 50 \text{ ms}$) and then proceed to the acquisition of the next shot.

7.3.4 Summary of general hypotheses on the acquisition model

In general, the current capabilities of SNAKE are currently restricted by the following hypotheses.

- (i) Tissue parameters and physiology are frozen during the acquisition of a single shot: T_1, T_2^*, ρ, χ are constant when the signal (7.5) is computed.
- (ii) The BOLD effect is linearly sensitive to TE.
- (iii) The BOLD effect does not modify the phase of the signal.
- (iv) Complex valued Gaussian noise is added in the k-space across coils (modeling the thermal noise). The noise level is calculated from a user-prescribed SNR.
- (v) The acquisition of shots follows Eq. (7.5). As § 7.6 will demonstrate, this model can be simplified to (7.6) if the trajectories do not have long readouts and are temporally smooth.

These are the minimal hypotheses on which the implementation of SNAKE is based. However, it also provides a flexible framework for adding more complex models (e.g. motion or other physiological perturbations) through the handler mechanism as described in § 7.4.1. To make SNAKE user-friendly and reduce its parameterization, additional hypotheses can be formulated. However, the potential bias introduced by these simplifications should be carefully considered. In § 7.5.1 we describe further restricting hypotheses for our study case.

It should be outlined that SNAKE does not aim to reproduce the full complexity of MR physics, but instead to provide a realistic and reproducible framework that allows us to simulate various acquisition scenarios in fMRI with full control over the ground truth and its parameterization. As such, SNAKE provides an upper limit on the statistical sensitivity/specificity compromise given the definition of an acquisition and image reconstruction set-up. Hence, SNAKE is defined as an instrumental framework for benchmarking fMRI reconstruction methods that comes with reconstruction methods and statistical analysis tools to perform end-to-end validation of fMRI acquisition and reconstruction methods, as depicted in Figure 7.3.

7.4 Main characteristics of SNAKE implementation

SNAKE has been designed as a fully reproducible modular fMRI simulator capable of generating k-space data efficiently. In what follows, we give a broad overview of the main features of SNAKE. Then they will be illustrated in § 7.5 and § 7.6.

7.4.1 Modular Approach

SNAKE adopts a modular approach to simulate 3D + time fMRI data.

²https://mind-inria.github.io/snake-fmri



Figure 7.3: Modular design of the SNAKE simulator, which embeds an ecosystem of Python packages for its different building blocks. 1. The dynamics of the acquisition are simulated, along with the sampling pattern for each shots. 2. The acquisition engine gathers the multiple shots involved in the acquisition as well as the current state of the 3D anatomical phantom. It relies on the MRI-nufft package. 3. Then the output k-space data is reconstructed into the image domain using some specific method implemented in pysap-fmri/ModOpt[Far+20], thereby producing an estimation of the fMRI time-series. 4. Finally, the sequence of fMRI volumes is analyzed through a general linear model (GLM) defined from the experimental paradigm and implemented in nilearn[Abr+14]. Statistical maps are produced and compared to the ground truth used at the simulation stage using confusion matrices and statistical metrics (accuracy or ACC, balanced accuracy or BACC).

Typically, the simulation begins with the definition of an anatomical phantom of the brain in the image domain (see Figure B.1 in § B.1), and proceeds to add BOLD contrast and various noise sources through object called "handlers" before generating the k-space data from a user-defined sampling pattern. As shown in Figure 7.3 those handlers can be chained to produce complex behaviors from simple operations. Moreover, numerous sampling patterns (especially non-Cartesian ones) can be generated through the use of the MRI-NUFFT library. The acquisition process is depicted in Figure 7.2. SNAKE supports both (7.5) and (7.6) through dedicated "engines" that are optimized for parallel computing of shots on GPU.

This option allows users to assess the need for a more complex and possibly more computationally demanding model in the context of the chosen scenario.

Moreover, SNAKE directly provides access to variational reconstruction methods using PySAP-fMRI³ and ModOpt [Far+20].

At the end of the processing chain, it is possible to compare the reconstructed fMRI images and time series data with ground-truth simulation, and evaluate the effect of acquisition parameters (SNR, acceleration factor, etc.) and reconstruction strategies (density compensated adjoint Fourier, compressed sensing reconstruction, etc.) on image quality, as well as on statistical sensitivity/specificity compromise. Statistical analysis is performed using the nilearn package [Abr+14].

7.4.2 Performance, reproducibility and scalability made easier for neuroscientists

Performance

Storing in RAM the full high spatiotemporal simulation is challenging at the high spatial and temporal resolution (340 Gigabytes are required for a 1mm-iso volume, with a unitary TR of 50ms shot-wise for 5 minutes of a typical fMRI run) and each substantial change, like adding noise would create a new copy of this data. Instead, we propose to yield the data to be acquired shot-wise on the fly, as each time point in the time series can be computed from a sequence of transformations applied to a single anatomical volume (for instance, adding the BOLD contrast, noise, or using motion parameters). Moreover, whenever possible, the computations are performed on GPU and shot-wise acquisitions are performed in parallel, the data is also eagerly off-loaded to hard disk when it is no longer required for computations. The k-space data generated by SNAKE is exported in the standardized ISMRMRD format (.mrd, inati_ismrm_2016).

Reproducibility

Enabling reproducibility in the study of fMRI processing methods and their benchmarking is at the heart of the development of SNAKE. This simulator can be installed directly from the Python package archive (https://pypi.org/project/snake-fmri/) and its core only depends on standard and well-tested Python packages. Simulation setup can be shared through .yaml files describing the recipe for building a simulation⁴.

³https://github.com/paquiteau/pysap-fmri

⁴The detail configuration and run scripts for scenarios of § 7.5 are available at https://github.com/ paquiteau/snake-fmri

Scalability and interoperatibility

Using the .mrd file as output, we open the door for interoperability of SNAKE with other toolboxes for image reconstruction, such as SigPy [OL19] or BART [UOT15]. We also provide optimized data-loader for the .mrd files generated by SNAKE.

Furthermore, based on the hydra framework [Yad19] we can run multiple simulations with different parameters or handlers and to perform image reconstruction of fMRI volumes and simple statistical data analysis to compare competing approaches. It also allows us to scale up simulations from a single laptop to high-performance computing clusters.

7.5 Numerical experiments

In this section, taking advantage of the modularity and scalability of SNAKE, we demonstrate the use of a controlled simulation framework to explore the potential benefits and challenges of moving to higher resolution in space and time for fMRI experiments. Results and analysis are presented in the next section.

7.5.1 Acquisition scenarios

First, we detail three simulated scenarios of increasing complexity (from 3D Cartesian low spatial resolution to 3D non-Cartesian high temporal or spatial resolution) with basic reconstruction and statistical analysis pipeline. All scenarios simulate a five-minute-long run with full-brain coverage during a simple visual stimulation using a standard block-design paradigm at 7T, which alternates 20s-on and 20s-off periods. In response to visual stimuli, we induce a 2. 5% change in BOLD contrast (following Eq. (7.7)) in a region of interest (ROI) in the occipital cortex. This ROI is defined from a fuzzy segmentation of gray matter that intersects an ellipse located in the occipital cortex.

Furthermore, we ran both the T_2^* relaxation model (7.5) and the simplified Fourier model (7.6) for the acquisition, to determine in which settings a more complicated model is required.

Working hypotheses to simplify acquisition

In all scenarios, the acquisition model considers only the intrinsic phenomenon of MR physics and lets aside all other modeling steps (motion, physiological noise, off-resonance effects) that could be compensated using a more complicated reconstruction setup. In particular, in addition to the fundamental hypothesis in § 7.3.4, we assume that:

- Following Hypothesis (iv) in § 7.3.4, the coil covariance matrix Σ is set to identity, and we use a user-defined SNR level ($SNR_i = 1000$, determined using the calibration step described in § 7.3.3).
- There is neither motion nor physiological perturbation (aside from the BOLD signal): in the context of acquisition and reconstruction benchmarking, we assume that the motion could be compensated, even though this task may require in practice navigator echoes or dedicated hardware.

- Off-resonance effects are not considered: From this point of view, we put ourselves in an ideal situation with minimal static and dynamic ΔB_0 inhomogeneities. Furthermore, we recently showed that it can be compensated using field cameras, in addition to a static map ΔB_0 in the forward model at reconstruction time. (cf.[Amo+24a; Amo+23]) (7.5).
- Reducing the number of tissues: Limiting ourselves to the three cortical tissues (white matter (WM), gray matter (GM), cerebrospinal fluid (CSF)) does not hinder the comparison between ground-truth and reconstructed images or statistical analysis.

Such assumptions allow us to focus on the core components of SNAKE (See § 7.3.4), reduce the computational load, but nonetheless provide an upper-bound on the performances of the tested scenarios. Possible extensions can be added by implementing the appropriate handlers.

The three scenarios and their computational cost are summarized in Table 7.4. Details of all parameters used are available in the .yaml configuration files in the SNAKE source repositoryfootnote 4.

As the unitary TR_{shot} and the simulation time (5 minutes) are the same for all scenarios, each simulation has a budget of 6000 shots to acquire. Depending on the number of shots allocated per frame (for a specific acceleration factor AF), we end up with a variable number of k-space volumes across the three scenarios.

The simulation times reported in Table 7.1 underline the additional overhead of using (7.5) instead of (7.6) for the simulation. In Scenario 1 the simulation time is limited by I/O communications, whereas Scenario 2 and 3 are computationally bounded.

Table 7.4: Overview of simulated scenarios and their computation requirements $-L, N_s, n_{jobs}$ specify the number of simulated coils, the number of shots acquired to create a k-space volume and the number of concurrently simulated shots, respectively.

C.J	Б		CND	r	N 7			n_{jobs}	Time		
Setup	Res.	Readout	SNR_i	L	N_s	1_{obs}	IR_{vol}		Fourier (7.6)	With T_2^* (7.5)	
Scenario 1	3mm	EPI	1000	1	44	$25 \mathrm{ms}$	2.2s	6	$3 \min 54 \sec$	3min55sec	
Scenario 2	3mm	SoS	1000	8	14	$30 \mathrm{ms}$	0.7s	6	1 min 31 sec	5min10sec	
Scenario 3	1mm	SPARKLING	30	32	48	$25 \mathrm{ms}$	2.4s	3	1h34min20sec	3h58min40sec	

Hardware: CPU: Intel i9-13900H, RAM: 32Go, GPU: NVIDIA RTX 2000 Ada. Common sequences parameters: $TR_{shot} = 50$ ms, TE = 25ms, $FA = 12^{\circ}$.

Scenario S1: 3D fully sampled Cartesian readout

As a first validation example, we simulated the acquisition of 3D Echo Planar Imaging (EPI) data, as an implementation of a 3D Cartesian readout. The 3D EPI is segmented plane by plane in the k space (as in Poser et al. [Pos+10]). Each slice was fully sampled at a 3mm isotropic resolution (matrix size: $60 \times 71 \times 60$), with a $TR_{shot} = 50$ ms. The volume-wise temporal resolution is $TR_{vol} = 2.2$ s at Ernst flip angle (12°), leading to a SNR_i of 1000 calibrated on real acquisitions (cf. § 7.3.3). Since the data were collected at the Nyquist rate, we restricted ourselves to a single-coil acquisition. This simple configuration can be simulated on a standard laptop in a few minutes, a similar configuration acquired with POSSUM would have taken several hours, for fewer slices (see Drobnjak et al. [Dro+06], Table 2), due to its more mathematically involved model related to MR physics.

Scenario S2: 3D under-sampled stack of spirals (SoS) readout with VDS acceleration along the stacking axis

The second scenario explored the possibility of SNAKE for accelerated imaging based on Compressed Sensing (CS) for data acquisition and image reconstruction. The resolution and FOV remained the same as for Scenario S1, but the goal was to increase the temporal resolution. To do so, we performed an acceleration on the 2D plane (k_x, k_y) using an in-out spiral acquisition and second we implemented a spiral stack using a variable density sampling along k_z , i.e. the stacking dimension (see parameters in Table 7.4). This sampling evolved across scans: Around 10% of the center planes were constantly acquired, while we used an acceleration factor AF=4 on the outer planes, as shown in Figure 7.5a. We eventually collected 14 spiral shots per volume, reaching in general $TR_{vol} = 0.7s$ as volume temporal resolution (TR_{vol}) . To compensate for aliasing artifacts, we simulated a multicoil acquisition with L = 8 receiver coils. Using GPU-accelerated NUFFT, the simulation time of the acquisition at 3mm isotropic resolution took around 1.5 minutes of computation (Table 7.4) in the CPU-based NUFFT implementation, which is also available, but remains slower.

This approach has already been studied numerically and experimentally in [PHA17] with image reconstruction strategies that take advantage of this acceleration mechanism, such as low rank + sparse regularization [PHA17; LF19; OCS15; CGM18]. Here, we limit image reconstruction strategies to frame-wise approaches, described in § 7.5.2.



Scenario 2, with a center-out order for acquiring shots in the stacking.



+ 0.50 0.25 0.00 - 0.25 - 0.50 0.50

Figure 7.5: Non-Cartesian k-space sampling trajectories used in Scenarios 2 and 3.

Scenario S3: Fully 3D SPARKLING under-sampled readout

We finally simulated a third scenario using a fully 3D non-Cartesian sampling pattern, based on 3D SPARKLING acquisition [Cha+22; Amo+24c] over L = 32 coils at an isotropic resolution of 1mm and $TR_{vol} = 2.4$ s. 3D SPARKLING implements a variable density sampling according to a prescribed sampling density in k-space, while complying with the hardware constraints on gradient magnitude G_{max} and slew rate S_{max} . These values are user-defined and were set to $G_{\text{max}} = 40 \text{mT/m}$ and $S_{\text{max}} = 180 \text{T/m/s}$. The target density was a radially decaying distribution parameterized by a cut-off and decay parameters set to (C, D) = (0.25, 2), see [Amo+24c, Eq. (3)] for details. Fully 3D SPARKLING allows us to further accelerate the acquisition process compared to Scenario 2 to reach higher spatial resolution. As we move to higher spatial resolution compared to scenarios S1 and S2 (voxels are $27 \times$ smaller), we multiplied the SNR_i level by a 30-fold factor.

The parameters for this scenario are based on the experimental setup of [Amo+24c], which is described in § 7.5.2, and in Figure 7.6, the upper row. In contrast to Scenario 2, here we adopted the "scan and repeat mode" that consists of sampling the same k-space locations across consecutive frames. Time-varying 4D SPARKLING acquisition for fMRI is left for future work and currently beyond the scope of SNAKE.

7.5.2 Reconstruction Strategies for Scenario 2 & 3

For simplicity, as the primary focus here was on the simulation of k-space data, we restrict ourselves to a classical, frame-wise CS based image reconstruction with a standard sparsity-enforcing regularization term in the wavelet domain:

$$\widehat{\boldsymbol{x}}_{t} = \arg\min_{\boldsymbol{x}\in\mathbb{C}^{M}} \frac{1}{2} \sum_{\ell=1}^{L} \|\boldsymbol{\mathcal{F}}_{t}\boldsymbol{\mathcal{S}}_{\ell}\boldsymbol{x} - \boldsymbol{y}_{t,\ell}\|_{2}^{2} + \mu_{t}g(\boldsymbol{\Psi}\boldsymbol{x}).$$
(7.10)

This means that each frame t is reconstructed independently of the others, as we use the following notation:

- \mathcal{F}_t is the Fourier transform operator for the trajectory at time t.
- Ψ is a orthogonal wavelet transform such as sym-8.
- $\mu_t > 0$ is the regularization parameter for frame t.
- $g(\cdot)$ is the regularization function, here a standard ℓ_1 -norm is used.

The cost function to be minimized is convex but non-smooth. Its global minimizer can be found iteratively using a wide range of proximal gradient methods with possible acceleration schemes such as FISTA, POGM [BT09; KF18]. As illustrated in Figure 7.3, CS-based image reconstruction was performed using the PySAP library [Far+20], and the implementation of the POGM algorithm [KF18]. Specifically for fMRI we implemented a dedicated plugin called **pysap-fmri**. To reduce the number of free parameters, the regularization parameter for each frame μ_t was estimated using Stein's unbiased risk estimate (SURE) principle [DJ94], as detailed in § B.2.

The versatility of SNAKE allowed us to investigate two different acquisition strategies in the space k within scenario S2, the first based on *static* SoS and a second one associated with a *dynamic* SoS. In the static regime, a constant spiral stack was used for all frames, whereas in the dynamic regime, a spiral stack that varied over time was designed by picking up k_z plans randomly for each and every frame, as shown in Figure 7.5a. Then CS reconstruction of each frame, as described in Figure 7.6, was carried out according to three different mechanisms:

(i) in a *cold-start* manner, where each volume in frame t is reconstructed by solving (7.10) independently of previous frames $\{1, \ldots, t-1\}$, as illustrated in Figure 7.6[top].

- (ii) in a warm-start manner, where the volume reconstructed in frame t by solving (7.10) is used as initialization to reconstruct the following volume frame t + 1, as shown in Figure 7.6[center].
- (iii) in a *smart* manner, using an additional *refined* initialization, where the last volume reconstructed in frame T using the previous warm-start strategy was eventually injected as a new set-up for all previous frames, as explained in Figure 7.6[bottom]. Coupled with dynamic SoS at acquisition, this approach is instrumental in visiting all k-space measurements across the consecutive frames.

All these variations are based on a frame-wise 3D image reconstruction strategy and, as such, provide a memory-efficient implementation. However, they do not leverage all 4D fMRI k-space data at once, in contrast to low-rank+sparse methods [PHA17; OCS15].

Estimation of μ_t in each frame (detailed in § B.2) plays a critical role in the performance of the *warm* and *refined* strategies. The first reconstructed frames are highly regularized, but as we progress towards the end of the run, the estimates of μ_t using the SURE-based methods get smaller, as we progressively embed more information to reconstruct volume x_t .

7.5.3 Evaluation methods

To evaluate the performance of each scenario and the impact of data acquisition and image reconstruction strategies, each combination was submitted to a standard fMRI statistical analysis pipeline, which consists of applying a general linear model (GLM) and testing the positivity of the single modeled experimental condition (visual) in the design matrix. Then a T statistic associated with the regression parameter β was formed and thresholded at p < 0.001 (one-sided), uncorrected for multiple comparisons.

As we know the ground-truth activated ROI, we can determine which detected activations are the true/false positives and negatives. The region of interest being small, we have a strongly imbalanced dataset and used the Precision/Recall curve instead of the classical receiver operator characteristic (ROC) curve to accurately compute for each scenario both the area-under-curve (AUC) and balanced accuracy (BACC) scores at p < 0.001. Additionally, we also measured image quality metrics (PSNR and SSIM) as well as temporal SNR (tSNR), to assess the quality of the reconstruction in space and time.

7.6 Results

7.6.1 Scenario S1

In Scenario 1, the k-space was fully sampled using EPI planes and reconstructed using the inverse FFT. The low spatial and temporal resolution results in high SNR and good image quality, as shown in Figure 7.7a. T_2^* relaxation only introduces negligible artifacts (5% of maximal error, see Figure 7.7b), and does not affect statistical performance compared to the basic Fourier model, as illustrated in Figure 7.7c. In general, this scenario validates the ability of SNAKE to handle end-to-end fMRI simulation and reconstruction.



Figure 7.6: Different methodologies for sequential reconstruction used in Scenario S2. Top: *cold start* reconstruction, each frame is reconstructed independently. Center: *warm start* reconstruction, each frame is reconstructed using the previous frame as initialization. Bottom: *refined* reconstruction, after a warm start reconstruction, the last frame is used as initialization for all other frames.



(a) Reconstructed fMRI volume for Scenario S1 for T_2^* model (7.4) (left) and Basic Fourier model (7.6) (right). The ground-truth ROI where activation lies is outlined in cyan. Detected activations surviving at p < 0.001, uncorrected, thresholding are overlaid using a colorjet map.



(b) Error between the T_2^* and Basic Fourier models, using the first reconstructed volume of each time series. The error has been rescaled by the maximum intensity of the basic Fourier model.



(c) Precision/Recall curves for S1 for the two models.

Figure 7.7: Comparison of the two acquisition models (basic Fourier and T_2^* effects) for scenario S1.

7.6.2 Scenario S2

Scenario S2 focuses on the joint effect of optimizing acquisition and reconstruction strategies for acceleration purposes and, therefore, demonstrates the versatility of SNAKE. As described in § 7.5.1 two acquisition strategies (static vs. dynamic spiral stack) compete and three variations of reconstruction methods (cold vs warm vs refined initialization for each volume).

The image quality and statistical maps are shown in Figure 7.8 and the extensive quantitative statistical analysis is reported in Figure 7.9. Based on these results we can make the following claims.



(b) Basic Fourier model



Figure 7.8: Activations maps for scenario S2. Top: T_2^* relaxation is taken into account. Bottom: Basic Fourier model, no T_2^* relaxation. Colored frames within each insert follow the convention adopted in the tables reported in Figure 7.9. Detected activations surviving at p < 0.001, uncorrected, thresholding are overlaid using a colorjet map.

(a) T_2^* model

SSIM PSNR Setup tSNR AUC BACC Label Acq. Recon First Last First Last dynamic COLD 0.640.5119.94 19.80 6.210.01 0.55dynamic REFINED 0.670.69 21.08 21.0323.640.340.80 dynamic WARM 0.640.79 19.94 21.06 12.950.01 0.58static COLD 0.640.64 19.91 19.91 43.08 0.70 0.96 REFINED static 0.60 0.60 20.4620.46 37.99 0.75 0.95 WARM 20.26static 0.640.7419.9120.180.350.84





(a) T_2^* model

(b) Basic Fourier model

	SSIM		PS	NR	+SND	AUC	PACC		
Label	Acq.	Recon.	First	Last	First	Last	1.5111	AUC	BACC
	dynamic	COLD	0.62	0.50	22.92	22.85	9.42	0.06	0.60
——	dynamic	REFINED	0.72	0.74	27.47	27.20	33.30	0.78	0.97
	dynamic	WARM	0.62	0.85	22.92	27.19	20.00	0.16	0.73
	static	COLD	0.62	0.62	22.95	22.96	71.01	<u>0.79</u>	0.97
	static	REFINED	0.65	0.66	25.13	25.14	49.36	0.88	0.99
	static	WARM	0.62	<u>0.78</u>	22.95	24.51	28.77	0.63	0.94

Figure 7.9: Quantitative metrics summarizing image quality (SSIM/PSNR), signal quality (tSNR) and statistical performances (AUC/BACC) for scenario S2. Top: T_2^* relaxation is taken into account. Bottom: Basic Fourier model. TSNR values are averaged over the occipital ROI. In the tables, First and Last refer to the first and last frames in the fMRI volumes sequences.

T_2^\ast relaxation must be taken into account for the long readout in non-Cartesian acquisitions

The T_2^* relaxation model must not be neglected in the simulation stage when considering non-Cartesian readouts instead of 3D EPI for comparable TE and observation time T_{obs} , as its impact on both image quality and statistical performance is significant. Notably, the comparison of results shown in Figure 7.9a and Figure 7.9b depicts a significant drop in SSIM, PSNR and tSNR scores for the T_2^* relaxation model due to blurring and contrast loss. This decrease in image quality happens because the reference MR image is the ideal phantom with the contrast at TE, which does not suffer from the T_2^* relaxation. Similarly, we observe a decrease in AUC/BACC for the T_2^* relaxation model, which is due to a systematic lower precision (higher false positive rate). This loss in precision is the result of lower z scores for that model compared to the basic Fourier model, as shown when contrasting Figure 7.8a and Figure 7.8b.

Hereafter, we only discuss the T_2^* relaxation model for this scenario, as it is the closest to real fMRI acquisition.

The refined initialization is key in dynamic acquisition for improved performances

When comparing image reconstruction strategies, we observe in Figure 7.9 that the warmstart mechanism is beneficial to improve the image quality between the first and last frames for both models. The gain in both SSIM and PSNR scores is actually large when this strategy is turned on. However, activating only the first warm start strategy does not yield improved statistical performance. Instead, using the *dynamic refined* approach in the reconstruction stage adds value to image quality and statistical performance, notably for the dynamic acquisition setup. In this specific scenario, this reconstruction strategy allows each frame to actually bring new information from complementary under-sampled k-space data and thus maintain good performances while increasing temporal resolution.

Moreover, the use of time-varying under-sampling patterns introduces temporal incoherence across frames and thus introduce different aliasing artifacts for each frame, which can be detrimental to image quality and statistical performances (see for instance the large gaps in the different scores between the *dynamic cold* and *refined* strategies in Figure 7.9).

Image quality is not a right proxy for good statistical performances

The static acquisition strategy provides the best statistical results, as it allows us to detect evoked brain activity in the targeted ROI with the highest precision/recall and tSNR scores (cf. Figure 7.9a). Here, the tSNR metric correlates well with AUC/BACC scores as the noise is purely thermal (no physiological noise). However, image quality is degraded in these cases; this is notably visible in Figure 7.8 when comparing the results *static cold* to *dynamic refined*. Strong aliasing and blurring artifacts actually corrupt reconstructed MR images in the static regime. Additionally, consistency of image quality (even in the degraded case) over the whole fMRI run also matters, as warm-start reconstruction (where image quality improves over time) shows poor statistical performances.

Finding the best trade-off between image quality and statistical performance can be achieved using the refined mechanism in terms of initialization strategy, as it provides the best image quality throughout the sequence of fMRI volumes while retaining good statistical performance. This is particularly noticeable in the dynamic acquisition setting: The AUC/BACC scores associated with the refined strategy (solid green curve) are higher than those associated with warm/cold strategies (solid blue and orange curves); see Figure 7.9a for the T_2^* model. More strikingly, this also happens in the static acquisition setting. A significant gain in SSIM/PSNR is observed with the refined strategy, and still the AUC/BACC scores are comparable to those measured using the cold initialization. When comparing the green and blue Precision / Recall dashed curves in Figure 7.9a, one can even observe that the green bullet is on top of the blue one, indicating a slight boost in precision (statistical sensitivity) at a given specificity level (same value along the y axis). Finally, all these observations that were established on the T_2^* model remain valid in the simplified Fourier model (cf. Figure 7.9b), which enforces these statements.

However, dynamic acquisition strategies have shown some great potential for increasing temporal resolution while maintaining good image quality. Preserving fine details in fMRI images at the output of the reconstruction pipeline matters for correct preprocessing of fMRI volumes and reliable statistical analysis (registration to a template, extraction of the cortical surface, etc.).

7.6.3 Scenario S3

Last, we present the results obtained for Scenario S3 (1mm isotropic resolution, $TR_{vol} = 2.4$ s) as a proof of concept of the scalability of SNAKE to high spatial resolution non-Cartesian fMRI.

(a) Activation maps. Left: Fourier model. Right: T_2^* model. (b) Precision Recall Curve.



(c) Quantitative results for Scenario 3.

Setup		SSIM		PS	NR	tSNR	AUC	BACC
Label	Acq.	First	Last	First	Last		лос	DAUU
	Basic Fourier	0.306	0.309	22.204	22.220	36.224	0.298	0.806
	T_2^*	<u>0.305</u>	<u>0.307</u>	<u>22.100</u>	$\underline{22.115}$	<u>34.894</u>	<u>0.229</u>	<u>0.774</u>

Figure 7.10: Results for scenario S3. Top left/center: First fMRI reconstructed volume using the standard Fourier (left) and T_2^* (center) models. Sequential cold-start CS reconstruction was performed frame-wise. Detected activations surviving at p < 0.001 thresholding are overlaid using a colorjet map. Top right: Precision/Recall curves for the two models. Accurate modeling of realistic T_2^* decay slightly impacts statistical performances. Bottom: Quantitative metrics summarizing image quality (SSIM/PSNR), signal quality (tSNR) and statistical performances (AUC/BACC).

With full GPU acceleration, sequential cold start reconstruction took 4min4sec per frame to converge, compared to 15-20min in the previous implementation [Amo+24c], thus requiring 8h30min for a complete fMRI run (125 frames). Both models (with and without T_2^* relaxation) show similar and relatively poor image quality due to blurring, as shown in Figure 7.10a. Additionally, the statistical performances are slightly lower (loss in sensitivity and specificity) when the T_2^* relaxation effect is considered in the data simulation process, as illustrated in Figure 7.10c.

7.7 Discussion

7.7.1 Effect of T_2^* relaxation

All scenarios have roughly the same T_2^* , TE and T_{obs} , but the T_2^* relaxation effect depends on the acquisition strategy and the k-space trajectory used. In Scenario 1 (EPI trajectory) no noticeable changes are visible, as the k-space is fully sampled and nearby k-space points are acquired at close times. In Scenario 2, the spiral trajectory allows for the acquisition of two neighboring points in k-space at both temporal extremities of the spiral. This results in a change of imaging contrast between two neighboring measurements in the k-space, thus inducing a loss of contrast in the reconstructed fMRI images. This effect can be alleviated using temporally smooth trajectories such as cones or MORE-SPARKLING [Gil+23] readouts. Statistical performance is also affected as lower z scores were retrieved for the T_2^* model compared to the basic Fourier model, as shown in Figure 7.8. However, for well-chosen reconstruction scenarios (e.g. static cold/refined), evoked activity is still well detected in the occipital ROI, and the precision-recall at p < 0.001 remains similar. Similar findings were observed in the case of dynamic acquisition with refined reconstruction.

In general, we recommend keeping the modeling of the relaxation effects T_2^* in the simulation despite the higher computing cost, as they can still influence the extraction of z scores as a function of the readout of k-space for given TE and T_{obs} .

7.7.2 Setting the noise level and tSNR

Concerning the choice of SNR_i , we reported in § B.3 the tSNR maps for each scenario. In S1 we measured a tSNR of 40.5 for the basic Fourier and T_2^* models in the occipital ROI. The tSNR increased to 120 in the CSF as shown in Figure B.2a. This is in line with what we could expect under experimental conditions. TSNR values for S2 are reported in Figure 7.9. As the T_2^* relaxation is significant in this scenario, we can see a drop of tSNR, which goes along with lower statistical performances. However, tSNR is not an oracle for statistical performance, as outlined in [Jam+21]. Finally, in S3, we get a tSNR of 36 in the occipital ROI. By decreasing SNR_i at higher resolution from 1000 at 3mm iso to 30 at 1mm iso, we obtained similar tSNR ranges in the three scenarios. Note that for S3, similar values have been empirically observed in [Amo+22a].

7.7.3 Insights from the different scenarios

The three scenarios presented in § 7.5 and 7.6 provide complementary views on the comparison of acquisition and reconstruction strategies.

Scenario S1 covers a low spatial and temporal resolution setting, uses standard Cartesian acquisition and basic reconstruction methods, and shows the potential of SNAKE to provide a basic reference for detecting evoked brain activity at a given input SNR. Beyond providing a simple and fast validation case for the simulator, this scenario could be used to compare the statistical performances of different competing experimental paradigms regarding the number and duration of stimuli, block vs. event-related designs, for a given scan time budget and under various artifacts.

Scenario S2 explores a low spatial, high temporal resolution setup and provides a large panorama of possible acquisition and reconstruction strategies. It shows the potential of SNAKE to produce a benchmark of competing techniques both at acquisition (static vs dynamic stack of spirals) and image reconstruction stages, particularly in the CS setting. This benchmark could be extended in the near future by using more complex reconstruction methods such as low-rank + sparse methods or even deep learning approaches. Interestingly, from its current implementation, we observed that the best image quality is not always associated with the highest statistical scores, and that the refined initialization strategy is key in dynamic acquisition for improved statistical performances.

Finally, Scenario S3 provides a high spatial and low temporal resolution setup, reaching the limits of current fMRI acquisition strategies for whole brain coverage. Unlike S2, the SPARKLING under-sampling pattern does not show aliasing artifacts, even at high spatial resolution, however, at the cost of some blurring. More generally, SNAKE offers new possibilities to optimize further 3D and even 4D under-sampling patterns.

In general, these three scenarios should be considered as an upper bound in terms of image quality and statistical performances for prospective validation on real fMRI experiments, as actual fMRI data acquisition and image reconstruction face additional issues such as imperfection in coil sensitivity estimation, presence of motion artifacts, off-resonance effects due to static and dynamic B_0 inhomogeneities.

7.7.4 Limits to the study

Our work is well defined by the hypothesis formulated in § 7.3.4, and their practical application in § 7.5. In particular, the scenarios presented in this chapter omit three major sources of noise in the fMRI data: head movement, physiological effects, and inhomogeneities of B_0 , since we focus on the core capabilities of SNAKE. Similarly, we only used the three main tissue classes (WM, GM, CSF). This allowed us to speed up the acquisition. However this choice was also the consequence of missing information in the literature about relaxation parameter values for other tissue classes at 7T. Moreover, the Cartesian scenario restricts itself to fully sampled EPI and lacks any acceleration setup such as GRAPPA or CAIPI [Bre+05; Gri+02]. Future work will address these aspects and explore more challenging reconstruction setups. However, SNAKE provides neuroscientists with an upper limit on the statistical impact and performance of many combinations of acquisition and reconstruction strategies, which offers an easy exploration of new methodologies.

Why off-resonance effects are not considered in this study

At first sight, the off-resonance effects due to static and dynamic B_0 inhomogeneities could be considered as a major source of artifacts in fMRI, especially for non-Cartesian acquisition strategy. However, even in the absence of their simulation, we can already see that it is already possible to discriminate acquisition and reconstruction strategies regarding their statistical performances in the downstream analysis. Furthermore, in reallife acquisition various methods are put in place to mitigates these effects (e.g. passive and active shimming, eddy current compensation, field monitoring, etc.). From this standpoint, the generated data using SNAKE could be considered a best-case scenario, that it is already possible to reach in experimental settings. We are comforted in this interpretation by the comparable image quality obtained by SNAKE (if not slightly worse) and experimental fMRI data, for instance in Scenario 3 which follows the experimental setup of Amor et al. [Amo+24c], where field camera was used to correct for static and dynamic off-resonance effects, as well as first order errors in the gradients waveforms. Finally, our experiments simulate activation in the primal visual cortex, which stays away from brain area (like ear canals and sinuses) where off-resonance effects are typically witnessed.

On a more practical side, modeling off-resonance effects faithfully would requires a dramatic increase in computational requirements, that would be detrimental to the ease of use of SNAKE, and the speed of the development cycle of new acquisition strategies. Modeling proper off-resonance effect would requires: (i) A high-resolution static field inhomogeneities map: To get such map, it would be need to increase the acquisition resolution of the simulation (e.g. from 3 mm^3 to 0.5 mm^3 and the number of tissues considered (from $N_{tis} = 3$ to $N_{tis} = 12$ in (7.5)), as they are all required to get a correct model of the air-tissues interfaces of the head. (ii) Applying such high-resolution field-map with high fidelity using (7.5) would also increase the number of interpolators to P = 100.

Thus, the number of Fourier Operators call being $N_{tis} \cdot P \cdot L$, it would leads to 1200 calls to high-resolution NUFFT operators instead of 400 lower resolution ones, and very high memory requirements with arrays of size $N_{tis} \times N_{voxels}$ for the phantom, $L \times N_{voxels}$ coil-sensitivity maps as well as a $P \times N_{samples}$ and $P \times N_{voxels}$ for the off-resonance maps. (resulting in $\simeq 70$ GB of GPU memory required per shot) In comparison, our Scenario 3 at 1 mm³ resolution looks far easier to compute.

Finally, without considering oracle knowledge of the field map and reconstruction at the target 3mm-iso will also be a challenge (with the increase in computation cost being proportional to the number of components used to decompose the field map). Overall, this would severely hinder the scalability of SNAKE and its ability to explore new strategies for both acquisition and reconstruction.

Restricting to single echo GRE fMRI

In this study, we focused on single echo GRE fMRI, which is the most common sequence used in fMRI experiments, in particular at ultra-high field. The aim of SNAKE is, in the first place, to push the boundary of what is possible in terms of spatial and temporal resolution for whole brain acquisition with ultra high field MRI. Thus, even if the modularity of SNAKE acquisition engine would allow it, we did not consider multi-echo GRE fMRI setup (A crude solution would be to run SNAKE for each echo and then combine the results, while ensuring that the overall parametrization is consistent across echoes and realistic). Furthermore, multi-echo setup are typically reduced in term of spatial resolution at ultra field, as multiple scans of the k-space needs to be perfomed during a shorter T_2^* decay.

Similarly, we did not consider spin-echo based sequence, they are not as common as GRE in fMRI and the extra acquisition time and SAR constraints make them less attractive for ultra-high field fMRI. However, the modularity of SNAKE acquisition engine would allow it to simulate such sequences, by replacing the GRE steady state signal model (7.2) with the appropriate spin-echo signal model.

7.7.5 Extending the simulator

As SNAKE is an open source software, external contributors from the fMRI community are welcome to participate in its extension to help refine the forward modeling to handle multiple sources of artifacts that contaminate actual fMRI data. In addition to head motion and off-resonance (of which we are adding some preliminary support), one may think of modeling of temporal aliasing artifacts due to physiological rhythms (heart beat or breathing rate) that are not sampled fast enough. Adequate fMRI acquisition and reconstruction methods could be studied with SNAKE to mitigate these additional sources of disturbance.

Similarly, more complex brain activation patterns spread over multiple ROIs (for using functional atlases from Yeo et al. [Yeo+11]) could be designed with spatial variations in the HRF shape, following the seminal work of PyHRF [Vin+14], or using meta-analyses and Python tools like neurosynth⁵ [yarkoni2011large] or Neuro-query⁶[dockes2020neuroquery] to define well-located activation peaks for given cognitive paradigms and tasks. Moreover, the BOLD effect could also be modeled to have an impact on the phase of the complex-valued signal[menon_functional_1992; KO12]. As the overwhelming majority of fMRI studies are conducted on the magnitude of the signal we did not consider these effects in this version of SNAKE.

So far, SNAKE has focused on brain mapping tasks in task-related fMRI. However, it could address the simulation and analysis of resting-state fMRI and be tuned to optimize the retrieval of resting-state networks from synthetic rs-fMRI data sets. In that regard, coupling with other fMRI simulators, such as Virtual Brain [Sch+22], could be instrumental in generating realistic rs-fMRI data as input reference data to SNAKE. As SNAKE is already interfaced with the nilearn package for statistical analysis and because the latter allows functional connectivity analysis, this extension to rs-fMRI data analysis could be quite straightforward.

However, the complexity of the simulation should be balanced with the need for the explicability of the results. On the one hand, adding layers (and their potential heavy parametrization) to the simulation will also induce a potential loss in explicability for the effect on downstream applications. On the other hand, the modularity of SNAKE allows the ablation study to enable / disable any aspect of the simulation that is to be analyzed which has the most significant impact.

7.7.6 Exploring the effect of tuning acquisition and reconstruction together

The first results obtained with SNAKE show the criticality of matching and tuning the experimental design, acquisition, and reconstruction strategies to obtain the best quality in the subsequent analysis. More complex acquisition and reconstruction methods could leverage the temporal redundancy in k-space data with global *a priori* such as low rank + sparse [PHA17; GMC22]. The *refined* strategy introduced in Scenario S2 shows some great potential for both static and dynamic acquisition schemes, and will be at the core of future development.

To the best of our knowledge, as summarized in Table 7.1, SNAKE is the only open source fMRI data simulator that can efficiently provide arbitrary fMRI k-space data. It is also the only one that can be used to benchmark reconstruction methods in an automated manner.

⁵https://neurosynth.org/.

 $^{^{6}}$ https://neuroquery.org/

7.7.7 More than a fMRI simulation tool

The modular approach of SNAKE also enables applications other than simple simulations. First, the combination of the running scheduling and analysis modules provides a reliable benchmark for image reconstruction methods (even for a single anatomical volume).

Second, each handler can also be viewed (and used) as a data augmentation layer for supervised deep learning methods, opening up new opportunities for fMRI image reconstruction. Currently, such models cannot be trained in the supervised setting due to the lack of ground truth (i.e. non-accelerated) real high resolution fMRI data. Alternative self-supervised approaches based on domain undersampling [demirel2021improved] are limited in terms of performance, notably in the high-resolution setting. Hence, SNAKE could be used to train deep neural networks dedicated to fMRI image reconstruction on synthetic fMRI data. These models might be fine-tuned later with transfer learning on real data sets. Alternatively, SNAKE might serve as a data augmentation tool. The recent review by Gopinath et al. [Gop+24] outlined the need for efficient synthetic data generation, a task for which SNAKE was precisely designed.

7.8 Conclusion

In this chapter, we have proposed a new fMRI data simulator, called SNAKE, which is packaged as an open source Python software, offering the ISMRM and OHBM communities the opportunity to advance the field of optimal fMRI data acquisition and image reconstruction at low scanning cost. More specifically, SNAKE is purposely designed to assess the impact of massively undersampled 3D non-Cartesian readouts aiming to reach unprecedented spatial or temporal resolution while maintaining whole-brain coverage, good image quality, and the ability to detect tiny BOLD effect through statistical guarantees. Through its modular design, SNAKE has been thought to remain open to external contributions as there is room to model additional aspects, notably external sources of artifacts (static and dynamic B_0 inhomogeneities, motion, etc.).

SNAKE comes also with an end-to-end and extensible reconstruction and statistical analysis pipeline. It provides the user with tools to reach new frontiers in fMRI data acquisition and image reconstruction strategies and to evaluate multiple competing scenarios from an image quality and statistical analysis viewpoint that cannot easily be ranked in advance. Future work will be dedicated to the interface SNAKE with deep learning for fMRI image reconstruction.

They don't give a damn about any trumpet playin' band It ain't what they call rock'n'roll — Sultans of Swing, Dire Straits

Chapter 8

Going deeper with Plug and Play Priors

Outline

Con	ntext and summary	121
8.1	Introduction	122
8.2	Plug-and-Play algorithms for MRI	123
	8.2.1 MR image reconstruction	123
	8.2.2 Plug-and-play algorithms	123
	8.2.3 Complex-valued denoisers for MRI	124
8.3	Proposed approach	124
	8.3.1 Denoised dataset generation	124
	8.3.2 Unsupervised preprocessing of multicoil data	125
	8.3.3 Proposed algorithm	125
	8.3.4 Preconditioning matrices	126
8.4	Experimental results and discussion	127
	8.4.1 Simulated non-Cartesian acquisition	127
	8.4.2 Training of the denoising prior	127
	8.4.3 Image reconstruction results	127
8.5	Conclusion & future work	129

The content of this chapter has been accepted for publication in the ISBI 2025 conference:

P.-A. Comby, B. Lapostolle, M. Terris, and P. Ciuciu. *Robust Plug-and-Play Methods for Highly Accelerated Non-Cartesian MRI Reconstruction*. Version 2. Oct. 2024. URL: https://hal.science/hal-04759015. Accepted to ISBI 2025

Context and summary

In Chapter 4 (§ 4.2.3) we introduced the general framework of Plug-and-Play (PnP) algorithms for image reconstruction. In this chapter, we focus on the specific application of MRI reconstruction, where the ill-posed nature of the inverse problem is particularly challenging. We propose a novel approach to MRI reconstruction by leveraging an annealed preconditioned Half-Quadratic Splitting (HQS) algorithm within the PnP framework. To avoid the limitations posed by the absence of high-quality, noise-free training data, we propose an unsupervised preprocessing pipeline that effectively denoises multicoil k-space data. This preprocessing involves a virtual coil combination followed by training an

unsupervised neural network for complex MRI signal denoising, generating a clean dataset suitable for subsequent training of a denoiser. By training a high-performance denoiser on these clean data and integrating them into various PnP algorithms, we demonstrate the superior reconstruction quality of our approach compared to existing methods.

This work is thus a first step in the development of PnP methods for fMRI reconstruction, that we expect to validate with SNAKE-generate data (Chapter 7) in the future.

8.1 Introduction

THE quest for improved resolution and shorter acquisition times in magnetic resonance imaging (MRI) has led to the development of more advanced acquisition strategies (see Chapter 2 for an overview) In this context, accelerated multicoil acquisition has become a standard practice in clinical settings, allowing faster scans while maintaining good image quality. However, as acceleration factors increase, less k-space samples are collected and image reconstruction becomes increasingly challenging due to the ill-posed nature of the underlying inverse problem.

To address this issue, reconstruction techniques based on Compressed Sensing (CS) theory have been widely adopted [LDP07]. Despite their success and the rise of a new generation of variable density undersampling patterns [Laz+19; Cha+22], CS-based methods often deliver suboptimal image quality, particularly at higher acceleration factors, where the ill-posedness of the problem is most severe. Recently, driven by the greater availability of public MRI data sets and computational power, deep learning methods have propelled image reconstruction quality to new heights [Zbo+18; Ham+18; Muc+21; Ram+22; PZJ23]. However, their performance is highly dependent on the acquisition setup (*e.g.*image resolution, sampling trajectory, acceleration factor, SNR of the acquisition, etc.) and are prone to overfitting and hallucinations.

Therefore, hybrid methods have emerged, which combine the capabilities of both the deep learning approach and the robustness of traditional CS methods. In particular, the Plug-and-Play (PnP) framework [VBW13; HLP22; Pes+21] has been shown to be a powerful approach to image restoration tasks and has been successfully applied to MRI reconstruction in various settings [Wei+21a; Ahm+20; Hon+24a; Fat+22]. However, the performance of PnP algorithms is strongly dependent on the quality of the learned prior, and training with appropriate data is key. This is problematic in applications dealing with indirectly acquired signals, such as in MRI, where a reconstruction algorithm is necessary to generate the reference data, potentially introducing biases, artifacts, or inaccuracies into the training data. Furthermore, the convergence properties of the PnP algorithm can be challenging to ensure [Hon+24a; Ter+24].

In this work, we propose to bypass both these issues by (i) denoising the multicoil fastMRI dataset in a fully unsupervised way and, (ii) applying an annealed half-quadratic-splitting (HQS) algorithm yielding fast reconstruction with state-of-the-art performance, while guaranteeing stable convergence.

8.2 Plug-and-Play algorithms for MRI

8.2.1 MR image reconstruction

In an ideal (i.e. artifact-free) setting, MR imaging consists of collecting the Fourier transform of the spatial distribution of nuclear magnetization in the organ of interest. In a multicoil (say L) acquisition setup, the signal measured by each coil ℓ surrounding the organ can be modelled as: for $\ell \in \{1, \ldots, L\}$,

$$y_{\ell} = \mathcal{F}_{\Omega} \mathcal{S}_{\ell} x + e_{\ell}, \tag{8.1}$$

where $x \in \mathbb{C}^N$ is the image to reconstruct, \mathcal{F}_{Ω} is the (potentially non-uniform) Fourier transform operator depending on a sampling trajectory Ω , $(\mathcal{S}_{\ell})_{1 \leq \ell \leq L}$ are the complexvalued spatial sensitivity maps (or S-maps), and e_{ℓ} is the realization of some random noise. The S maps are smooth low-frequency images and can be estimated from the central frequencies of y_{ℓ} . The goal of the reconstruction is then to recover the image xfrom the data $(y_{\ell})_{1 \leq \ell \leq L}$.

8.2.2 Plug-and-play algorithms

Traditional image reconstruction algorithms propose to solve the ill-posed problem (8.1) by reformulating it as a minimization problem of the form

$$\widehat{x} = \arg\min_{x \in \mathbb{C}^N} \frac{1}{2} \sum_{\ell=1}^L \|\mathcal{F}_{\Omega} \mathcal{S}_{\ell} x - y_\ell\|^2 + \lambda r(x),$$
(8.2)

where r is a convex function and $\lambda > 0$ a regularization parameter. In this context, the choice of an appropriate regularizer r is paramount to ensure good reconstruction quality to counterbalance the ill-posedness of (8.1). Denoting $f(x) = \frac{1}{2} \sum_{\ell=1}^{L} ||\mathcal{F}_{\Omega} \mathcal{S}_{\ell} x - y_{\ell}||^2$, the solution \hat{x} is then obtained by iterating a proximal algorithm [CP11a], a classical instance being the proximal gradient algorithm

$$x_{k+1} = \operatorname{prox}_{\gamma\lambda r}(x_k - \gamma \nabla f(x_k)), \qquad (8.3)$$

where we recall that $\operatorname{prox}_{\gamma\lambda r}(x) \stackrel{\text{def}}{=} \operatorname{argmin}_{u} \gamma\lambda r(u) + \frac{1}{2}||x - u||_2^2$. Plug-and-Play (PnP) algorithms [VBW13; Pes+21; HLP22] propose to replace the proximal operator in (8.3) with a denoising neural network D_{σ} , where $\sigma > 0$ is the denoising power. There, D_{σ} accounts for an implicit prior more expressive than standard regularizers r, and the noise level σ plays a similar role as λ . PnP algorithms have been shown to strongly outperform their traditional counterparts in image restoration [Zha+21b], MRI [Hon+24a], and astronomical imaging [Dab+22] to name a few.

PnP algorithms are owing their popularity to several advantages. Firstly, the sophisticated implicit prior D_{σ} is trained independently from the acquisition strategy in(8.1): If the under-sampling pattern Ω changes, the same D_{σ} can be used for the reconstruction, which constitutes a significant advantage over unrolled algorithms. Secondly, while the resulting algorithms may be unstable, annealing schemes [Zha+21b] allow us to stabilize the algorithm while drastically improving the reconstruction. As a consequence, PnP algorithms have been used for MRI reconstruction in a number of works [CJ23; Hon+24a; Wei+21a]. Much of these works outline the importance of using convergent reconstruction algorithms. In this work, we are rather leveraging reconstruction algorithms with annealing strategies.

8.2.3 Complex-valued denoisers for MRI

At the algorithmic level, the proximal operator in (8.3) takes complex variables as input. Similarly, the denoiser used in the PnP algorithm will be faced with complex images with real and imaginary parts once a multicoil combination has been performed at each gradient step. Therefore, PnP algorithms for multicoil MRI data need to rely on denoisers that handle complex data.

Moreover, the quality of the reconstruction of the PnP algorithms is highly dependent on that of the denoiser at hand [Ter+23]. However, in the MRI context, since the image is indirectly acquired and often requires preprocessing, either linear in the parallel imaging setting or nonlinear, such as (8.3), no data set of ground truths (noiseless) is available, and only noisy datasets (e.g. fastMRI) are available. Several works have proposed to denoise MR images [CLY22; Pfa+23], but to our knowledge, none applied it to PnP algorithms. In this work, we denoise the fastMRI dataset in an unsupervised way to train our denoisers.

8.3 Proposed approach

In this section, we detail our proposed approach. We first introduce our strategy for producing a clean dataset of complex-valued images from multicoil MR acquisitions, which will serve as a basis for training our denoising DNN. We next introduce the proposed preconditioned HQS algorithm and the choices considered for the preconditioning matrix.

8.3.1 Denoised dataset generation

We rely on the fastMRI dataset [Zbo+18] for our study, which provides only fully sampled multicoil data. To reconstruct a single complex image from multicoil measurements, we propose a virtual coil combination method. This approach transforms each coil's k-space data into the image domain using FFT, aligns the phase information, and combines the coil images into a single complex image x_c [Par+14]. We apply this operation to the full multicoil training dataset of fastMRI { $(y_1^1, \ldots, y_L^1), \ldots, (y_1^N, \ldots, y_L^N)$ }, yielding a new dataset of N complex images { x_c^1, \ldots, x_c^N }.

The dataset still contains residual acquisition noise, which can be significant in some cases. To address this issue, we propose using these complex, noisy images to train a restoration neural network, f_{θ} , in an unsupervised manner with the Neighbor2Neighbor loss [Hua+21]. This approach allows us to preprocess the data without relying on clean, noiseless samples. Specifically, for a given noisy sample x_c the training loss is computed as follows:

$$\mathcal{L}(x) = \|f_{\theta}(g_1(x_c)) - g_2(x_c)\|_2^2$$

$$+ \eta \|f_{\theta}(g_1(x_c)) - g_2(x_c) - (g_1(f_{\theta}(x_c)) - g_2(f_{\theta}(x_c)))\|_2^2,$$
(8.4)



Figure 8.1: Effect of the unsupervised preprocessing on the training dataset. Top row shows a non-preprocessed sample of the fastMRI multicoil dataset, obtained with virtual coil combination procedure described in [Par+14]; bottom row shows the preprocessed image.

where g_1 and g_2 are sub-sampling operations and $\eta > 0$ a regularization parameter.

After training, we apply the learned model to the full dataset of combined images $\{x_c^1, \ldots, x_c^N\}$, yielding a new training dataset of N denoised, complex images $\mathcal{D}_{\text{train}} = \{x^1, \ldots, x^N\}$ that will serve as a basis for training our implicit prior D_{σ} .

8.3.2 Unsupervised preprocessing of multicoil data

We apply the procedure described in Section 8.3.1 to the fastMRI brain multicoil dataset. We choose a DRUNet model [Zha+21b] for f_{θ} in (8.4), with an input and output convolutions channel number set to 2 for real and imaginary parts, without conditioning on the noise level. Training is performed with a batch of size 10 and on randomly extracted patches of size 128×128. We use the Adam optimizer with learning rate 10⁻⁴, and set $\eta = 2.0$ in (8.4). We then preprocess the full training dataset with f_{θ} . An illustration of the resulting preprocessed dataset is given in Fig. 8.1. Notice that the resulting image shows much lower residual noise, while maintaining a high resolution.

8.3.3 Proposed algorithm

We propose to investigate several variants of PnP algorithms. We first study the standard PnP version of the PGD algorithm (8.3) where the proximal operator is replaced by a denoising neural network D, writing

$$u_k = x_k - \gamma P \nabla f(x_k)$$

$$x_{k+1} = \mathcal{D}_{\sigma}(u_k),$$
(PnP-PGD)

where P is a preconditioning matrix, and $\gamma > 0$ and $\sigma > 0$ are stepsizes and noise levels, respectively. While the use of P = Id remains predominant in the image restoration literature, recent works have shown that using carefully tailored preconditioners strongly improves the reconstruction quality in MRI [Hon+24a].

Following the DPIR approach [Zha+21b], we propose to extend the preconditioning of [Hon+24a] and use a half-quadratic splitting algorithm coupled with an annealing strategy summarized hereafter:

$$u_{k} = \operatorname{prox}_{\gamma_{k}f}^{P}(x_{k})$$

$$x_{k+1} = D_{\sigma_{k}}(u_{k}),$$
(PnP-HQS)

where $\operatorname{prox}_{\gamma f}^{P}$ is the proximity operator of γf in the metric induced by P, and $(\gamma_k)_{1 \leq k \leq K}$ (resp. $(\sigma_k)_{1 \leq k \leq K}$) are sequences of stepsizes (resp. noise levels). We underline that, in the case of $P = \operatorname{Id}$, (PnP-HQS) reduces to the DPIR algorithm from [Zha+21b]. Importantly, we point out that despite their similarities, (PnP-PGD) and (PnP-HQS) differ fundamentally, as illustrated by the following result: Assume that there exists a convex function g such that $D_{\sigma} = \operatorname{prox}_{g}$. Furthermore, assume that $P = \operatorname{Id}$ and that for all k, $\gamma_k = \gamma$ and $\sigma_k = \sigma$ in (PnP-HQS). Then:

- 1. (PnP-PGD) converges to $x^* = \underset{\sim}{\operatorname{argmin}} \gamma f(x) + g(x)$ [Pes+21].
- 2. (PnP-HQS) converges to $x^* = \underset{x}{\operatorname{argmin}} \gamma f(x) + {}^1g(x)$, where 1g denotes the Moreau envelope [Bau+17] of q.

We note that the two algorithms converge to different points. Although the minimization problem associated with (PnP-PGD) seems more natural than that of (PnP-HQS), empirical results show that the latter tends to outperform the former. Secondly, the assumption that $D = \text{prox}_g$ for some convex function g is well-known to not hold in practice, resulting in non-convergent and potentially unstable algorithms.

Therefore, numerous works have proposed to enforce the convergence of the algorithm through Lipschitz constraints [Pes+21; HLP22] or through preconditioning strategies [Hon+24a]. Inspired by [Zha+21b], we propose to set $(\gamma_k)_{1 \leq k \leq K}$ and $(\sigma_k)_{1 \leq k \leq K}$ as decreasing sequences according to the rules $\sigma_k = \sigma_0 \xi^{\delta}$ and $\gamma_k = \lambda \sigma_k$, where $\delta, \lambda > 0$ and $0 < \xi < 1$ are tunable hyperparameters. This can be interpreted as a way to force convergence of $(x_k)_{1 \leq k \leq K}$, thereby circumventing the instability issues witnessed with (PnP-PGD).

8.3.4 Preconditioning matrices

In [Hon+24a], the authors show that introducing preconditioning improves both the stability of the algorithm and the quality of the reconstruction. We propose to extend their approach to the case of the HQS algorithm (PnP-HQS). In our case, the preconditioning is applied to the metrics of the proximal operator of the data-fidelity term $\operatorname{prox}_{\gamma f}^{P}(x) \stackrel{\text{def}}{=} \operatorname{argmin} \gamma f(u) + \frac{1}{2} ||x - u||_{P}^{2}$, where we recall that, for all x, $||x||_{P}^{2} = \langle x, Px \rangle$. In contrast to (8.3) where the preconditioning applies straightforwardly to the gradient, computing $\operatorname{prox}_{\gamma f}^{P}(x)$ is more involved. In our case, we rely on a subiteration solver.

Following [Hon+24a], we investigate three choices for matrix P. Denoting $A \stackrel{\text{def}}{=} \mathcal{F}_{\Omega} \otimes \mathcal{S}_{\ell}$ we consider the "F-1" preconditioner $P = 2 - \alpha A^{\mathcal{H}}A$, and the Chebychev preconditioner $P = 4 - \frac{10}{3}A^{\mathcal{H}}A$. In a nutshell, the choice of these preconditioners arises from finding matrices P minimizing the spectral radius of Id $-\alpha PA^{\mathcal{H}}A$, see [Hon+24a] for more details. Lastly, we also investigate the case P = Id corresponding to the standard PnP approach.



Figure 8.2: Top row: Reconstruction results for (8.1) for HQS and PNP algorithms with no (Id) or static (F1) Preconditioner at AF=4 (top) and AF=16 (bottom). PSNR/SSIM metrics are shown inside each image. Bottom right inset depicts the residual maps ($5 \times$ magnified).

8.4 Experimental results and discussion

8.4.1 Simulated non-Cartesian acquisition

Our ground-truth data consists of raw, fully-sampled multi-coil brain k-space data from the FastMRI dataset. To simulate an accelerated acquisition, we used a spiral pattern with either AF=4 or AF=8 (resp., 4- or 8-fold under-sampling). On the resulting under-sampled k-space, we added white Gaussian white noise $(e_l \text{ in } (8.1))$ with variance $\nu_l = 10^{-4} \cdot \max(\sum_{\ell} |x_{\ell}|^2)$. The spatial sensitivity maps $(S_{\ell})_{1 \leq \ell \leq L}$ in (8.1) were estimated from the center of the k-space (20×20 Hamming window), transformed in the image domain, and we masked-out the resulting images to fit the brain.

8.4.2 Training of the denoising prior

We trained our denoising prior D_{σ} on the preprocessed data set $\mathcal{D}_{\text{train}}$ at hand from Section 8.3.1. Following standard practice from the PnP literature, we chose a DRUNet model [Zha+21b] conditioned on the noise level σ and with input and output convolutions of the model containing 2 channels accounting for real and imaginary parts. Additionally, the model was trained in a supervised way with an ℓ_1 loss on the proposed preprocessed dataset. Training was performed for 100 epochs for noise levels σ ranging in $[0, 5 \cdot 10^{-5}]$, with the Adam optimizer with learning rate 10^{-4} , with batch of size 10 and on randomly extracted patches of size 128×128 . In contrast to other work in the MRI community where denoisers were trained on normalized images in [0, 1], we proceeded differently as our denoisers worked on unnormalized MRI slices.

8.4.3 Image reconstruction results

Preconditioning boosts image quality

To evaluate the quality of our reconstructions, we use the Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) as our primary metrics. Both metrics are computed on the magnitude of the output images, and the PSNR is restricted to

Table 8.3: Reconstruction metrics for different solvers and acceleration factors. Reported values are an average of 5 trials. **First** and <u>second</u> best performances are highlighted per AF (column-wise), across all PnP-like methods. Hyperparameters were chosen/NCPDNET was trained at AF=4.

ec.	Solvor	\mathbf{PS}	NR	SSIM		
$\mathbf{P}_{\mathbf{I}}$	Solver	AF=4	AF=16	AF=4	AF=16	
leb	HQS	34.521	33.251	0.917	0.902	
C	PNP	33.025	28.223	0.908	0.812	
Ţ	HQS	35.874	33.641	0.921	0.911	
μīų	PNP	34.596	30.280	<u>0.932</u>	0.865	
PN	IP-FISTA	27.493	29.657	0.791	0.808	
PN	IP-Dyn	32.695	28.111	0.903	0.809	
q	HQS	35.245	30.437	0.937	0.905	
Π	PNP	34.043	32.118	0.892	0.874	
$ \mathbf{A} $	FISTA-Wavelet	33.230	29.753	0.923	0.867	
Z	NCPDNET	37.765	29.912	0.977	0.916	

the brain region of interest to avoid bias from the background. We show in Fig. 8.2 reconstruction results for AF=4 and AF=16 for PnP reconstructions. As in [Hon+24a] we observe that preconditioning significantly improves image quality compared to non-accelerated counterparts (HQS-Id and PnP-Id). The visual improvements are in line with the broader quantitative assessment in Tab. 8.3. Moreover, the HQS splitting scheme improves reconstruction over the standard PGD and FISTA PnP algorithms.

Generalization to other acceleration factors

The fully unrolled NCPDNET approach [Ram+22] still outperforms the PnP and HQS methods, when trained for a specific AF. Yet as the AF increases, the PnP methods shine, as they can maintain good image quality, without requiring retraining or fine-tuning a neural network. At AF=16 (Tab. 8.3), all PnP methods outperform NCPDNet and variational approaches, without further fine-tuning (for each method, the same hyperparameters were used across all AF).

Stability of PnP methods and further acceleration

We show in Fig. 8.4 the convergence plots for all PnP and HQS algorithms at AF=4 Interestingly, we notice that the HQS algorithm is slower than its PGD counterpart with preconditioning, which indicates that further fine-tuning of the hyperparameters could lead to faster results. In particular, the stepsizes could be adaptively learned on a validation set. Concerning the preconditioned PnP (PnP-F1, PnP-Cheb, PnP-Dyn)[Hon+24a], we



Figure 8.4: Evolution of the PSNR across iterations for various PnP algorithms at AF=16. Dashed lines correspond to variations of the (PnP-PGD) algorithm, while solid lines correspond to variations of the (PnP-HQS) algorithm.

did use a normalization equivariant network. Surprisingly, we were unable to reproduce the convergent behavior for those methods as shown in Fig. 8.4. In contrast, we relied on the annealing scheme for the HQS iterations to improve image quality.

8.5 Conclusion & future work

In this work, we introduced a novel algorithm for multicoil, non-Cartesian MRI reconstruction by leveraging an annealed preconditioned Half-Quadratic Splitting (HQS) algorithm within the Plug-and-Play (PnP) framework. To avoid the limitations posed by the absence of high-quality, noise-free training data, we proposed an unsupervised preprocessing pipeline that effectively denoises multicoil k-space data. This preprocessing involves a virtual coil combination followed by training an unsupervised neural network for complex MRI signal denoising, generating a clean dataset suitable for subsequent training of a denoiser.

By training a high-performance denoiser on these clean data and integrating it into various PnP algorithms, we demonstrated the superior reconstruction quality of our approach compared to existing methods. Further work includes extension to 3D and 2D+time data, such as those considered in cardiac imaging or functional MRI. This work is thus a first step toward the development of PnP methods for fMRI reconstruction, that we expect to validate using SNAKE-generated artifical data (Chapter 7) in the near future.

Detailed settings, implementations, and a reproducible benchmark are available at https://github.com/paquiteau/benchmark_mri_pnp. The benchmark was realized using the Benchopt framework [Mor+22] and the deepinv¹ and MRI-NUFFT² libraries.

¹https://github.com/deepinv/deepinv

²https://github.com/mind-inria/mri-nufft

It's a dangerous business, Frodo, going out your door. You step onto the road, and if you don't keep your feet, there's no knowing where you might be swept off to. — Bilbo Baggins, The lord of the rings

Chapter 9

Plug-and-Play reconstruction for 3d non-Cartesian fMRI data

Outline

Context and summary	. 131
9.1 Introduction \ldots	. 132
9.2 Materials and Methods	. 132
9.2.1 Acquisition simulation setup	. 132
9.2.2 Reconstruction with PnP prior	. 134
9.2.3 Neural network architecture	. 135
9.3 Experimental results	. 136
9.3.1 Benchmark of MR image reconstruction methods	. 136
9.3.2 Training	. 136
9.3.3 MR image reconstruction results	. 136
9.3.4 Statistical analysis of fMRI sequences	. 137
9.4 Conclusion and Perspectives	. 138

The content of this chapter has been submitted for publication at EUSIPCO 2025 conference:

Pierre-Antoine Comby, Matthieu Terris, Alexandre Vignaud, and Philippe Ciuciu. *Plug-and-Play Reconstruction for 3D Non-Cartesian fMRI Data*. Mar. 2025. URL: https://hal.science/hal-04993532. Submitted to EUSIPCO 2025

Context and summary

In this chapter, we combine the advances made in Chapter 7 and Chapter 8. Notably, we use the SNAKE simulator to evaluated the image quality and statistical performances of the preconditioned HQS algorithm for fMRI use. In a benchmark against compressed sensing strategies, we show that PnP-based reconstruction strategies produce high-quality images but suffer from a loss in statistical sensitivity compared to compressed sensing reconstructions

9.1 Introduction

In fMRI, as time becomes an additional dimension in data acquisition and image reconstruction, strategies must be carefully tailored to balance spatial resolution, temporal resolution, signal-to-noise ratio (SNR), and brain coverage. High SNR can be secured by performing scans at ultra-high magnetic field (7Tesla or above). However, the current race towards higher spatial and temporal resolution, ultimately to differentiate brain activity patterns across cortical layers (e.g., at 500µm resolution) [Pol+18; BHF21] or to uncover fast hemodynamic responses (e.g., at 0.5 Hz) [VP21], remains an unmet challenge for a complete brain coverage. An additional difficulty is the substantial variability of fMRI responses (evoked brain activity) across individuals, which impedes the optimization of this acquisition/reconstruction pipeline in vivo in healthy volunteers as a ground truth is missing.

We propose optimizing the data acquisition/image reconstruction pipeline *in silico* using the SNAKE simulator [CVC24b] to overcome this issue. SNAKE allows us to synthesize realistic artificial fMRI data under various acquisition setups and, hence, to select the best according to different metrics, namely image quality and the statistical sensitivity/specificity compromise, which are not necessarily aligned. In this work, we introduce a novel learning-based prior, based on a 3D implicit denoiser, injected into a Plug-and-Play (PnP) method for fMRI image reconstruction and compare the performance of deep learning and variational methods for image reconstruction on both image quality and statistical criteria.

The remainder of this paper is structured as follows. First, we describe the simulation setup to produce realistic and accelerated artificial task-related fMRI data in k-space. Then we provide some background on PnP methods based on denoising priors and justify our choice of neural network architecture. Next, we introduce our benchmark that compares zero-filled adjoint Fourier, Compressed Sensing, and deep learning for fMRI image reconstruction on a frame-by-frame basis. Finally, we present and discuss our benchmark results by assessing each approach's ability to recover decent image quality and good statistical performances to detect evoked brain activity.

9.2 Materials and Methods

9.2.1 Acquisition simulation setup

To generate realistic artificial fMRI data in both the time-image and k-space domains, we use the SNAKE simulator [CVC24b]. We first consider a brain phantom, depicted in Fig. 9.1(a), with a true T_2^* -weighted contrast. To achieve this, we simulate a 3D Gradient Recall Echo acquisition (repetition time TR = 50 ms, echo time TE = 22 ms, flip angle $FA = 12^\circ$) using an 8-channel phased array coil (L = 8). Next, we simulate BOLD activity according to the block design paradigm (black trace in Fig.9.1(b)), alternating 20 seconds-long blocks of visual stimulation and rest in the primary visual cortex based on the Harvard-Oxford atlas. The brain region eliciting the evoked activity is delineated by the red contour in Fig.9.1(a). The magnitude of the BOLD effect is defined by multiplying the fuzzy gray matter tissue mask voxel-wise with the mask of the primary visual cortex. In each voxel, we then generate the simulated BOLD time series by scaling
the block-design regressor (black trace in Fig.9.1(b)) with the magnitude of the BOLD effect and convolving the result with a canonical hemodynamic response, resulting in the orange trace in Fig.9.1(b). The stimulation lasts T = 300 sec in total (one half shown in Fig.9.1b.



(a) Brain Phantom on a field of view of (192,192,144) mm with the T_2^* -weighted contrast prescribed by TR, TE, FA. Activated region in the occipital cortex in red.



(b) Example of a block design paradigm for fMRI.

Figure 9.1: Brain Phantom used in our simulation (top) and average BOLD response (simulated and measured) in the ROI (bottom).

To simulate a realistic MR acquisition process, we consider the presence of N_{tis} tissue classes in the 3D brain phantom. This means that we model tissue-specific parameters, namely tissue proportion $(w_i(\mathbf{r}))_{i=1}^{N_{\text{tis}}}$ per voxel \mathbf{r} , contrast $((\mu_i)_{i=1}^{N_{\text{tis}}})_{i=1}$ in the prescribed GRE sequence configuration, and transverse relaxation times $(T_{2,i}^*)_{i=1}^{N_{\text{tis}}}$ to accurately simulate a spatially varying MR imaging contrast, embodying T_2^* -decay. Then, regarding the simulation of fMRI data in k-space, we use a stack of in-out spirals as sampling trajectories $(\mathbf{k}_s(t))_{1\leq s\leq n_{\text{shots}}}$ for n_{shots} consecutive shots of duration TR, at each timestep $0 \leq t \leq T$. The center of each spiral is placed at TE to optimize the T_2^* contrast. Since we start with a standard 3-mm isotropic resolution, full k-space coverage at the Nyquist rate typically consists of 48 spirals, resulting in a volumetric TR (TR_v) of 2.4 sec. Here, as we are interested in high temporal resolution $(TR_v = 0.75 \text{ seconds})$, we subsample the k-space data by a factor of 3.2 along the stacking dimension k_z (see Fig. 9.2), which consists of retaining only $n_f = 15$ k-space planes (3 central planes constantly sampled and 12 randomly picked up). As a consequence, this procedure produces an acceleration factor of 3.2 and $T/TR_v = 400$ fMRI scans. In this context, the simulated data $y_{\ell,s}(t)$ for the s-th shot on the ℓ -th coil follows an extended Fourier model:

$$y_{\ell,s}(t) = \int_{FOV} \sum_{i=1}^{N_{\text{tis}}} w_i(\boldsymbol{r}) \, \mu_i \, e^{-t/T_{2,i}^*} \mathcal{S}_{\ell}(\boldsymbol{r}) e^{-2i\pi \boldsymbol{k}_s(t) \cdot \boldsymbol{r}} d\boldsymbol{r} + \varepsilon_{\ell,s}(t), \qquad (9.1)$$

where $\varepsilon_{\ell,s} \sim \mathcal{N}(0, \Sigma), \forall \ell = 1, \dots, L$ is the thermal noise with homoscedastic coil covariance matrix Σ defined as follows:

$$\boldsymbol{\Sigma} = \frac{1}{SNR} \left(\int_{FOV} \sum_{i=1}^{N_{\text{tis}}} w_i(\boldsymbol{r}) \mu_i d\boldsymbol{r} \right) \boldsymbol{I}_L, \qquad (9.2)$$

where SNR refers to the signal-to-noise ratio. Importantly, the BOLD activation is updated at every unitary TR, modulating the contrast μ_i within a designated Region of Interest (ROI), referred to as the Activation ROI tissue. This region's intensity varies over time according to the prescribed BOLD effect for a given shot t, with a modulation amplitude of 2.5% relative to the baseline contrast ($\mu_{ROI} = 0.025 \cdot h(t)$, where h(t) is the prescribed hemodynamic response. The maximum variation of $y_{\ell,s}(t)$ due to the BOLD effect is therefore impacted similarly. More details on the physics and computational aspects of the simulation process are available in [CVC24b].

Once all shots $y_t = (y_{\ell,s}(t))_{s,\ell}$ are sequentially generated according to Eq. (9.1), each group of n_f consecutive shots is combined to reconstruct a single volume, forming a frame in the fMRI sequence. It should be noted that information about the different tissue types or relaxation parameters is not available during the reconstruction stage. Eventually, the acquisition model for a single sought image x_t writes:

$$\boldsymbol{y}_{t} = \sum_{\ell=1}^{L} \mathcal{F}_{\Omega_{t}} \mathcal{S}_{\ell} \boldsymbol{x}_{t} + \boldsymbol{\varepsilon}_{t,\ell}$$
(9.3)

where $\Omega_t = \{ \mathbf{k}_s \mid tn_f \leq s \leq (t+1)n_f \}$ gathers all the shots of the *t*-th frame.

9.2.2 Reconstruction with PnP prior

Traditional methods for solving (9.3) often propose to estimate each x_t by solving a regularized minimization problem of the form

$$\widehat{\boldsymbol{x}}_t = \operatorname*{argmin}_{\boldsymbol{x} \in \mathbb{C}^N} f(\boldsymbol{x}, \boldsymbol{y}_t) + \lambda r(\boldsymbol{x}), \qquad (9.4)$$

where $f(\cdot, \boldsymbol{y}_t)$ is a data-fidelity enforcing functional (ensuring that the solution $\hat{\boldsymbol{x}}_t$ satisfies (9.3)), $r: \mathbb{C}^N \to \mathbb{R}$ ensures regularity of the solution, and $\lambda > 0$ is a regularization parameter. Here, f is chosen as $f(\boldsymbol{x}, \boldsymbol{y}_t) = \frac{1}{2} \sum_{\ell=1}^{L} ||F_{\Omega_t} \mathcal{S}_\ell \boldsymbol{x} - \boldsymbol{y}_{t,\ell}||_2^2$, while r is chosen to promote the sparsity of the reconstruction in an appropriate domain. The minimizer $\hat{\boldsymbol{x}}_t$ is then computed with an optimization algorithm, typically involving proximity operators and/or gradients of f and r [CP11b].

In this context, plug-and-play (PnP) algorithms propose to replace r in (9.3) with an implicitly defined regularization, typically through a deep neural network (DNN) trained as a denoiser. In practice, this operator replaces the proximity operator of r that appears in the minimization algorithm to solve (9.4). Following [Com+25b], we propose to perform the image reconstruction with a preconditioned Half-Quadratic Splitting (HQS)



Figure 9.2: Stack of spirals used to sample a single frame in 3D k-space. This pattern is repeated over time to collect the full fMRI sequence.

algorithm. More precisely, given a preconditioning matrix P and a number of iterations K, our reconstruction algorithm for estimating \hat{x}_t reads, for $0 \leq k \leq K$,

$$\begin{aligned} & \boldsymbol{u}_{t,k} = \mathrm{prox}_{\gamma f}^{P}(\boldsymbol{x}_{t,k}) \\ & \boldsymbol{x}_{t,k+1} = \mathrm{D}_{\sigma_{k}}(\boldsymbol{u}_{t,k}), \end{aligned} \tag{PnP-HQS}$$

where $\operatorname{prox}_{\gamma f}^{P}$ is the proximity operator of γf in the metric induced by P, $(\sigma_k)_{1 \leq k \leq K}$ is a decreasing sequence of noise levels, and where we set $\boldsymbol{x}_{t,0} = A^{\dagger} \boldsymbol{y}_{t}$. Interestingly, when $P = \operatorname{Id}$, (PnP-HQS) reduces to the DPIR algorithm from [Zha+21a] that has shown state-of-the-art performance in low-level vision tasks. In our experiments, we set $\gamma = 13$, $\sigma_k = \sigma_0 \xi^k$, with $\sigma_0 = 4 \cdot 10^{-3}$ and $\xi = 0.97$, and K = 20. Here, we consider the "F-1" preconditioner $P = 2 - \alpha A^{\mathcal{H}} A$ introduced in [Hon+24b],

Here, we consider the "F-1" preconditioner $P = 2 - \alpha A^{\mathcal{H}} A$ introduced in [Hon+24b], where $A = F_{\Omega_t} \otimes S_{\ell}$, and α is the Lipschitz constant of ∇f . Although other PnP algorithms may be chosen, two main observations motivate our choice: First, due to its annealing strategy of the parameters, it is fast and does not show the instabilities that other PnP algorithms may show [Com+24b]. Second, introducing the preconditioning matrix P ensures a robust reconstruction despite the high ill-conditioning of F_{Ω_t} in non uniformly sampled settings.

9.2.3 Neural network architecture

We next detail the chosen architecture for our denoiser D_{σ} . In this work, we consider complex volumetric data for $(\boldsymbol{x}_t)_{1 \leq t \leq T}$. More precisely, for all t, \boldsymbol{x}_t can be represented as a tensor of shape (2, D, H, W), where the first two dimensions account for real and imaginary parts of the data. Consequently, the denoiser to be considered should be able to handle complex volumetric images. We therefore propose to extend the 2D bias-free DRUNet architecture from [Zha+21a] into a 3D architecture by replacing the 2D convolution layers with the 3D convolutional layers, with 2 input channels. The resulting architecture enjoys the same theoretical properties as the original 2D architecture, *e.g.* its homogeneous property $D_{\alpha\sigma}(\alpha \boldsymbol{y}) = \alpha D_{\sigma}(\boldsymbol{y})$ for any $\alpha > 0$, ensuring good generalization abilities for this denoiser over images with various intensity ranges.

9.3 Experimental results



Figure 9.3: Z-score maps for CG/CS/HQS-F1 reconstruction methods (from left to right). the target ROI (voxel with at least 50% of gray matter, and belonging to primary visual cortex) is delineated in cyan. The complete anatomical ROI (occipital cortex) is contoured in green.

9.3.1 Benchmark of MR image reconstruction methods

We compare our reconstruction method with a wavelet-based compressed sensing reconstruction as implemented in pysap-mri [Far+20]. More precisely, we use (9.4) with $r(\boldsymbol{x}) = \|\Psi \boldsymbol{x}\|_1$, where Ψ is a wavelet decomposition. Eq. (9.4) can be efficiently solved with a standard forward-backward algorithm [CP11a]. We also provide comparisons with the pseudo-inverse $A^{\dagger}\boldsymbol{y}_t$, obtained by running a conjugate gradient (CG) algorithm on $f(\boldsymbol{x}, \boldsymbol{y}_t)$ for each frame t.

9.3.2 Training

We train our model on the Calgary-Campinas dataset [Sou+18]. We extract complex images by performing an inverse FFT on the raw Cartesian data in the k-space that will serve as our reference and combine multichannel inputs into a single 3D complex-valued volume using a virtual coil combination method [Par+14]. We train our model using the Adam optimizer with a 10^{-4} learning rate. Pre-training is performed for 200k iterations on randomly extracted patches of size $64 \times 64 \times 64$ with a batch size of 12. The model is then fine-tuned in larger patches $128 \times 128 \times 128$ with a batch size of 2 for 100k iterations.

9.3.3 MR image reconstruction results

We now present the main results of our study. First, we compare the image quality of reconstructed fMRI sequences using CG, CS, and PnP with the ground truth using PSNR and SSIM scores in Tab. 9.4. In Fig. 9.3, we present the three reconstruction methods considered. The background image of each case corresponds to the first volume reconstructed in the fMRI sequence, and the selected cuts maximize the number of voxels displayed for the target ROI. The learning-based PnP-HQS reconstruction provides the best image quality. In contrast, the wavelet-based CS reconstruction provides lower image

Table 9.4: Quantitative metrics on image quality and statistical performances. SSIM and PSNR scores are averaged over all frames. The balanced accuracy (BACC) is computed with a threshold set to $p < 10^{-3}$ and AUC and BACC were computed with respect to voxels with at least 50% of gray matter.

Name	PSNR	SSIM	tSNR	AUC	BACC
 CG	18.560	0.367	nan	0.782	0.754
 \mathbf{CS}	18.748	0.459	6.706	0.758	0.948
 HQS-F1	19.338	0.577	6.286	0.574	0.891

quality (0.12 lower SSIM scores) and embodies a few spiral aliasing artifacts. However, both regularized methods outperformed the classical CG reconstruction, which shows detrimental aliasing artifacts in the frontal and occipital regions.

9.3.4 Statistical analysis of fMRI sequences

Once reconstructed, we performed the regression analysis of fMRI signals using a simple general linear model (GLM) consisting of the regressor shown in orange in Fig. 9.1b and a constant baseline. The goal is to detect voxels eliciting evoked activity that correlates with the experimental paradigm. Then, z scores were calculated as the ratio of estimated regression weighted divided by the square root of the residuals in the GLM. Finally, z scores were thresholded at a p-value of p = 0.001 uncorrected for multiple comparisons to assess each reconstruction pipeline's statistical performances (sensitivity/specificity trade-off) and analyze to what extent the original ROI was recovered by computing scores and confusion metrics (true positive, false positive, etc.). As the ROI is small compared to the brain volume, we used precision and recall (true positive rate) to analyze the distribution of detected voxels.



Figure 9.5: Precision-Recall characteristic curve for the different thresholds of gray matter considered in the ROI.

Furthermore, voxels located in ROI have a spatially varying proportion of gray mat-

ter (GM), which directly impacts the magnitude of the BOLD effect in a proportional way. Thus, we considered different thresholds of z scores to define true positives. In Fig. 9.5, we present three different classification results for different GM proportions (voxels with respectively at least 10%, 50%, or 90%) in the respective dotted, dashed, and solid lines. Last, in Tab. 9.4, we compare the statistical performances for the CG/CS/HQS-F1 methods. As the CG method gives the highest AUC score while reporting the lowest image quality, this indicates a poor point spread function (PSF), which eventually blurs the true activations and produces false positives in addition to true positives. This is confirmed in Fig.9.5 by lower precision and higher recall scores for CG reconstruction compared to CS and HQS-F1. In contrast, CS reconstruction reports the best AUC/BACC and highest Precision/Recall scores despite aliasing artifacts and poorer image quality compared to HQS-F1. Hence, CS achieves the best compromise between image quality and statistical performance. The lower performance of the learning-based PnP method might be due to the fact that the 3D DRUNet has been trained on 1-mm isotropic images, while fMRI images were generated at 3-mm. Hence, additional fMRI simulations operating at 1-mm isotropic are necessary to confirm this finding soon, or instead to demonstrate the superiority of learning-based reconstruction methods, which have been trained at the target spatial resolution. However, because of their computational cost, these numerical experiments and the corresponding results were not available prior to submission.

9.4 Conclusion and Perspectives

In this work, we studied how a learning-based prior embedded in a PnP method impacts image reconstruction in 3D non-Cartesian high-temporal fMRI. To do so, we used the SNAKE-fMRI simulator to draw preliminary conclusions in a fully controlled environment. Although image quality is significantly improved using the learning-based PnP method compared to standard techniques (CG, CS), our findings suggest a compromise between image quality and statistical performance in terms of true positives (recall)/false positive (precision) balance that is currently optimized by CS reconstruction. This could result from training a denoiser on a dataset of T_1 weighted MR images, while the target application is fMRI and involves a T_2^* weighed contrast. Furthermore, the resolution change between the training set and the test set could also explain the loss in statistical performance. we observed with PnP.

Future work will therefore focus on simulating scenarios at matched spatial resolution (i.e., 1-mm isotropic) and possibly on matching the contrasts between the training and test stages. This extension is computationally and memory challenging.

For the sake of reproducible science, detailed settings, implementations, and reproducible benchmark are available at https://github.com/paquiteau/eusipco-pnp-fmri. The study was carried out using the deepinv¹ and SNAKE-fMRI ² libraries.

¹https://github.com/deepinv/deepinv

²https://github.com/mind-inria/snake-fmri

Tu seras Chef, veux tu ? Non pas aujourd'hui ni demain peut-être ; Mais lorsque l'heure sera venue: Lorsque ta vie sera droite comme une tige de roseau et simple comme un chant d'alouette — Aimée Degallier-Martin

Chapter 10 General conclusion

At the conclusion of these three and a half years of research, I am proud to have made meaningful contributions to the field of fMRI acquisition and reconstruction. With the invaluable support of my supervisors and in collaboration with the MIND team and partners in France and internationally, I developed open-source tools for the MRI community and established a solid foundation for advancing fMRI methodologies. This PhD thesis aimed to achieve higher spatio-temporal resolution in fMRI, currently offering a powerful simulation tool for studying brain function and unlocking new insights into neuroscience.

10.1 Contributions and limitations

The initial objective of this thesis was to develop novel learning-based reconstruction methods for fMRI data, aiming to enhance image quality and increase statistical power in task-related fMRI. However, we encountered several significant challenges along the way. First, the lack of high-quality, fully sampled ground truth raw fMRI datasets in the k-space posed a major obstacle, making it impossible to directly apply existing supervised learning-based methods. Second, alternatives relying on self-supervised data under-sampling, which have shown success in anatomical MRI, faced limitations in the context of fMRI. These methods struggle in the high-resolution regime (e.g. 1mm isotropic whole brain) due to the lack of sufficient k-space shots or trajectories. Accelerating 3D fMRI data acquisition actually exacerbates this issue by requiring more aggressive k-space undersampling, which, in turn, reduces the size of the training and test sets when splitting the data for learning and validation purposes.

To address the scarcity of reference data, we began by leveraging existing fMRI datasets to evaluate the performance of denoising methods such as NORDIC[Viz+21], which operate post-reconstruction. During this process, we discovered that other well-established statistical approaches could outperform NORDIC, offering a notable improvement in the trade-off between statistical sensitivity and specificity in task-related fMRI. This realization led to the development of the patch-denoise software (detailed in Chapter 6), which has garnered significant attention from the fMRI community and neuroscientist colleagues at Neurospin.

Nonetheless, we became firmly convinced of the critical importance of accessing large ground-truth raw k-space datasets in fMRI to effectively benchmark future advancements in learning-based reconstruction methods. Given the prohibitive costs of collecting extensive fMRI datasets at 7T and the challenges in obtaining fully sampled high-resolution fMRI k-space measurements, we opted to develop a novel simulator, named SNAKE—an

acronym for Simulator from Neuro-Activation to K-space Evaluation (described in Chapter 7).

SNAKE has proven invaluable in faithfully replicating fMRI acquisition scenarios, including complex 3D non-Cartesian imaging setups, through artificial simulations. It enables the generation of fully sampled high-resolution fMRI k-space data based on a given experimental design in task-related fMRI. Furthermore, SNAKE simulates the entire workflow of accelerated fMRI data acquisition and image reconstruction. To ensure the computational efficiency of simulating high-resolution non-Cartesian fMRI data, we heavily invested in the development of MRI-NUFFT, a specialized package that integrates multiple interfaces for the NUFFT transform and accounts for MRI-specific features, such as multicoil acquisitions (Chapter 5). I took the role of the core developer of this package while being supported by other colleagues.

With SNAKE, we now have the capability to provide comprehensive end-to-end validation of fMRI reconstruction methods. Our current efforts focus on integrating deep learning techniques into the fMRI reconstruction pipeline. The initial step in this direction was the development of robust Plug-and-Play methods for 2D MRI reconstruction, which we aim to extend to 3D and functional MRI in the near future. These extensions hold promise for significantly enhancing reconstruction accuracy and efficiency. However, for broader applicability and impact, further improvements to SNAKE are necessary. In particular, the current version does not account for B_0 inhomogeneities, an important factor that needs to be addressed in future developments.

Additionally, we evaluated the performance of sequential reconstruction methods (Chapter 3) using SNAKE, which allowed us to rigorously assess their capabilities in various simulated fMRI scenarios. Moving forward, we plan to benchmark other popular approaches, such as Low-Rank+Sparse methods, to further explore their potential in advancing fMRI reconstruction techniques.

10.2 Perspectives

While this PhD thesis has contributed valuable insights and solutions to the fMRI reconstruction problem, it has also unveiled new questions that warrant further investigation. These can be summarized as follows:

10.2.1 Denoising functional MRI data

Additional testing of local-low-rank denoising methods

The potential of local-low-rank denoising methods for fMRI preprocessing is evident, as demonstrated in Chapter 6. However, our current benchmark could be significantly extended to better evaluate the benefits and limitations of these techniques. First, the activation detection methodology assumes full rank to determine the degrees of freedom, which might not hold true after the patches are denoised and recombined. A more robust evaluation could involve applying local-low-rank denoising methods to downstream tasks such as brain decoding or classification of task performances. The IBC dataset, which provides medium-resolution task-related fMRI data (1.5 mm isotropic for sixty fMRI protocols) and has been extensively studied at MIND, offers a promising testbed for such investigations. Notably, preliminary work in this direction was conducted during the internship of Marie Doggo, which I supervised, and this could serve as a foundation for further research. Second, the current implementation of local-low-rank denoising is not optimized for computational efficiency. Accelerating this process could greatly enhance its usability in practice. Specifically, leveraging GPU-based computation for complex-valued singular value decomposition (SVD) and optimizing the patch extraction step—similar to a convolution operation—could significantly improve the speed and scalability of the method. By addressing these aspects, we could enhance the practicality and robustness of local-low-rank denoising methods for fMRI and expand their application in neuroimaging workflows.

Explore new denoising methods

In addition to local-low-rank denoising, deeper methods, such as diffusion-based models, hold significant promise for fMRI data denoising. Diffusion models have already demonstrated impressive results in various image denoising tasks, and their application to fMRI could offer new opportunities for improving data quality. By leveraging the power of these models, we may be able to more effectively remove noise while preserving critical features, particularly in the challenging context of fMRI where noise can obscure subtle brain activations.

Furthermore, developing robust denoisers is essential, especially for their potential integration into Plug-and-Play (PnP) frameworks. In such architectures, a learned denoiser could replace the traditional proximal operator, offering a more flexible and data-driven approach to regularization. By improving these denoising methods, we could enhance both the effectiveness and efficiency of PnP algorithms in fMRI reconstruction tasks. This line of research could lead to more effective, scalable, and robust solutions for fMRI preprocessing, making it an exciting area for further investigation.

10.2.2 fMRI Reconstruction

Perform a full benchmark of fMRI reconstruction methods

Chapter 3 outlined the broad spectrum of reconstruction strategies applicable to fMRI. In Chapter 7, we utilized SNAKE to evaluate sequential reconstruction methods with a warm-start strategy. However, this analysis could be expanded by incorporating a wider range of reconstruction techniques. A more comprehensive benchmark is needed, evaluating these methods using consistent performance metrics such as image quality and the statistical performance of downstream analysis (e.g., brain activity detection and decoding tasks). The ultimate goal is to identify the optimal acquisition-reconstruction pairing for specific fMRI paradigms, accounting for factors such as spatial and temporal resolution, noise levels, and data acquisition strategies. This ongoing effort will continue to drive the development of more efficient and accurate reconstruction methods tailored to the unique challenges of fMRI.

Data augmentation for deep learning with SNAKE

SNAKE is a powerful tool that can be leveraged to generate realistic artificial fMRI data. One exciting application would be using it to create a large dataset of 3D+t fMRI image sequences, which could then be utilized to train deep learning methods for fMRI reconstruction. By generating synthetic high-resolution, fully-sampled data from specific experimental designs, SNAKE enables the creation of an ideal supervised setting for deep learning model training. Such a dataset would provide the necessary ground truth for a variety of reconstruction methods, including plug-and-play approaches. Moreover, by varying acquisition parameters, noise levels, and temporal dynamics, this synthetic data could help train models that are both robust and capable of handling a range of fMRI acquisition protocols. This approach addresses the challenge of data scarcity and the need for high-quality supervised datasets, ultimately enabling the development of more powerful and generalizable fMRI reconstruction models.

Jointly optimizing acquisition and reconstruction for fMRI using SNAKE

The SNAKE simulator offers a promising avenue for jointly optimizing both the fMRI acquisition process and the subsequent image reconstruction. This could be achieved by leveraging SNAKE's ability to generate realistic artificial fMRI data, which could then be used to guide the optimization of both acquisition parameters (such as sampling patterns) and reconstruction techniques. This approach could build upon existing frameworks like PROJeCTOR [RC23], which was originally developed for anatomical imaging. By extending PROJeCTOR to the fMRI domain, SNAKE could be used to simulate a wide range of acquisition strategies and reconstruction methods, providing a comprehensive framework for optimization.

Improving and extending SNAKE

SNAKE has proven to be an invaluable tool for simulating and evaluating fMRI acquisition and reconstruction strategies. However, there are several potential improvements that could make it even more powerful and versatile:

- Support for additional pulse sequences: Currently, SNAKE is optimized for certain acquisition protocols (GRE sequence), but it could be extended to include support for other pulse sequences like Spin-Echo or Multi-Echo GRE. This would enable more diverse and realistic simulations of different fMRI acquisition schemes.
- Incorporating physiological noise and other variables: The modularity of SNAKE makes it well-suited to add additional model components, such as physiological noise (e.g., respiratory and cardiac fluctuations), variability in the hemodynamic response function (HRF), and resting state brain activity. Including these factors could lead to more realistic and comprehensive simulations of fMRI data, improving the fidelity of results and offering insights into how these factors influence reconstruction performance.
- Physiologically driven brain activity simulation: SNAKE could be further improved by integrating with advanced, physiologically-driven brain activity models, such as

The Virtual Brain[San+13]. This would allow for more biologically accurate and dynamic simulation of brain activity, providing more realistic data to assess the performance of various reconstruction algorithms.

• Efficiency enhancements: A key area for improvement is the efficiency of SNAKE. One potential optimization is the implementation of a caching mechanism for repeated shots within a simulation. By reusing previously computed results, the computational burden could be reduced, significantly improving simulation speed.

10.2.3 Deep-learning for fMRI reconstruction

Deep learning-based methods for fMRI reconstruction are still in their early stages. The size and dimensionality of fMRI data, combined with the lack of high-quality supervised datasets, present significant challenges for applying classical supervised learning approaches. However, the remarkable results achieved in anatomical MRI reconstruction using deep learning cannot be overlooked, and we firmly believe that deep learning holds great promise for fMRI reconstruction.

One promising approach to address the computational challenges in fMRI reconstruction is the use of Plug-and-Play (PnP) methods, which combine the strengths of deep learning with traditional iterative optimization algorithms. These methods are not only more efficient but also more flexible in handling different acquisition strategies, making them suitable for real-time applications in fMRI. We provided preliminary result in Chapter 9.

Even when using simulated data (e.g., generated by SNAKE) for training, there remains a concern about the potential discrepancy between the synthetic data and real experimental data. Factors such as subject-specific variability, noise sources, and acquisition-related distortions are difficult to perfectly replicate in simulations. This gap could impact the generalization ability of deep learning models, making them less effective when applied to real-world fMRI data.

10.2.4 Towards in-vivo utilization

One of the major challenges in advancing fMRI reconstruction methods is bridging the gap between simulation-based studies and real-world, in-vivo data. While SNAKE has proven to be a powerful tool for simulating artificial fMRI data, the ultimate goal is to validate these methods in real-world scenarios. It is essential to compare the insights gained from numerical experiments on synthetic data generated by SNAKE with actual in-vivo fMRI data to ensure the methods' reliability and effectiveness. To achieve this, we propose to explore transfer learning techniques based on domain adaptation. This approach involves adjusting the statistical distribution of synthetic fMRI data to align more closely with that of real fMRI data. By identifying and addressing any potential distribution shifts between synthetic and real datasets, we can improve the simulator, bringing SNAKE closer to producing more accurate and realistic data representations. The primary aim is to fine-tune the simulator to reflect the nuances and variability found in actual experimental data, especially across different MRI field strengths such as 3T, 7T, and even 11.7T. This will help ensure that the methods developed using SNAKE are not only theoretically sound but also applicable to practical, real-world fMRI data. If significant discrepancies are identified between synthetic and real data, they should be

Chapter 10 General conclusion

addressed by updating the simulator to account for these differences, leading to even more robust and accurate simulations. Through this iterative process, combining synthetic data with in-vivo validation, we will move closer to achieving high-fidelity fMRI reconstruction methods that can be applied in clinical and research settings.



Appendix A

Inverse problems and convex optimization essentials

Outline

A.1	Inverse Problems and Convex Optimization	147
	A.1.1 Introduction to inverse problem	147
	A.1.2 Convex analysis essentials	149
A.2	Optimization algorithms	151
	A.2.1 Gradient Descent	152
	A.2.2 Splitting methods	152
	A.2.3 Alternating Direction Method of Multipliers (ADMM)	154
	A.2.4 More than two functions	155

THIS chapter introduces the mathematical context and tools useful for the reconstruction of fMRI data presented in Chapter 2 and Chapter 3. However it remains generic, and the reader is encouraged to skip it if they are already familiar with the concepts of inverse problems, convex optimization and the use of deep learning in this context. The application of these tools to fMRI data will be detailed in the next chapter.

A.1 Inverse Problems and Convex Optimization

A.1.1 Introduction to inverse problem

In the fields of computational imaging and signal processing and physical measurements, a common mathematical framework is typically used. It typically consists of:

- \mathcal{H}, \mathcal{G} two Hilbert space of finite dimensions, such as \mathbb{R}^N or \mathbb{C}^N .
- $A: \mathcal{H} \to \mathcal{G}$ a linear operator representing the measurement of a quantity in \mathcal{G} from a sample represented in \mathcal{H} .
- Some error in this model, represented as noise $\varepsilon \in \mathcal{G}$.

Finding solutions \boldsymbol{x} of

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{\varepsilon} \tag{A.1}$$

is known as solving the inverse problem. Due to the presence of noise ε , and the properties of A, this is not necessarily an easy task.

Definition A.1 (Hadamar's Conditions). An inverse problem is said to be well-posed if it meets all the following conditions:

- *i. Existence: There exists a solution* \boldsymbol{x} *to* (A.1) *for every* \boldsymbol{y} *.*
- ii. Uniqueness: This solution is unique.
- iii. Stability: The solution \boldsymbol{x} depends continuously on \boldsymbol{y} (A small change of \boldsymbol{y} implies a small change of \boldsymbol{x}).

In any other case, the problem is said to be ill-posed.

A first approach to solve (A.1) is to use the least-square minimization

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}\in\mathbb{C}^N} \frac{1}{2} \|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}\|_2^2$$
(A.2)

$$= \mathbf{A}^{\dagger} \mathbf{y} \tag{A.3}$$

Where $\mathbf{A}^{\dagger} = (\mathbf{A}^{H}\mathbf{A})^{-1}\mathbf{A}^{H}$ is the Moore-Penrose pseudo-inverse of \mathbf{A} . Yet, this solution is most of the time not suitable in practice for the following reason:

- The solution might not be unique. The Moore-Penrose pseudo-inverse only returns the minimal norm solution (which may not be desired)
- If $\|A\hat{x} y\|_2^2 \neq 0$ then \hat{x} is not a solution for (A.1).
- If the condition number $\chi(\mathbf{A}) = \sigma_{\max}(\mathbf{A})/\sigma_{\min}(\mathbf{A})$ (*i.e.*, the ratio of the maximal and minimal non-zero singular values of \mathbf{A}) is not acceptable (typically in the order of 1–10) then even a small presence of noise $\boldsymbol{\varepsilon}$ can yield to solution that are far from the expected solution.
- Computing A^{\dagger} explicitly would require computing $(\mathbf{A}^{H}\mathbf{A})^{-1}$ which is not tractable for big dimensions. At least, a conjugated gradient method (see Theorem A.7) would be used instead.

To overcome the limitation of the least-square methods, the main approach is to add an extra term to (A.2) to enforce some extra properties on the estimate $\hat{\boldsymbol{x}}$. This is typically done by adding an extra penalty term $g(\boldsymbol{x})$:

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}\in\mathbb{C}^{N}} \underbrace{\frac{1}{2} \|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}\|_{2}^{2}}_{f(\boldsymbol{x})} + g(\boldsymbol{x})$$
(A.4)

 $f(\boldsymbol{x})$ is the *data-consistency* term, and encodes all information about measurements on \boldsymbol{x} , on the other hand, $g(\boldsymbol{x})$ is the *prior* and represents the knowledge and assumption we might have on \boldsymbol{x} directly.

Remark. In a probabilistic setting, the least-square solution arises as the maximum likelihood (ML) solution when the noise follows a white Gaussian probability. Introducing g then yields the maximum a posteriori (MAP) estimate and $g \propto \log \pi(x)$, where $\pi(x)$ is the prior distribution on x.

If sufficient assumptions are made on f and g (notably their convexity), the inverse problem will satisfy the Hadamar's conditions. Finding efficient algorithms for finding such minimum under various condition on f and g is an on-going quest in mathematics, whose main results useful for our case are summarized hereafter.¹

A.1.2 Convex analysis essentials

In this section, we will introduce the required assumptions and method that are used to solve Eq. (A.4). It is just the tip of the iceberg of a rich literature. The following is adapted from several books [SMF15; BV04; BC11] and articles[CP21; CP11a].

Following the usual notations, in this section the data will lie in a Hilbert space \mathcal{H} of finite dimension, and the measurement in \mathcal{G} , with the forward *linear* operator $\mathbf{A} : \mathcal{H} \to \mathcal{G}$ (meaning that A could be represented as a matrix).

Definition A.2. Let $f : \mathcal{H} \to \mathbb{R}$

(i) f is proper if dom $f = \{x | f(x) < \infty\} \neq \emptyset$.

(ii) f is lower semi-continuous if its epigraph is closed:

 $epi f = \{(x, t) \in \mathcal{H} \times \mathbb{R} | f(x) \leq t, x \in \text{dom } f \}.$

- (iii) f is convex if $\forall x, y \in \text{dom } f, \forall \alpha \in [0, 1] f(\alpha x + (1 \alpha)y) \leq \alpha f(x) + (1 \alpha)f(y)$.
- (iv) We note $\Gamma_0(\mathcal{H})$ the set of functions that verifies (i), (ii) and (iii).
- (v) f is β -smooth if it is differentiable and its gradient is Lipschitz continuous i.e.:

$$\|\nabla f(x) - \nabla f(y)\| \leq \beta \|x - y\| \quad \forall x, y \in \mathcal{H}$$
(A.5)

(vi) Let $f \in \Gamma_0(\mathcal{H})$, The subdifferential in $x \in \text{dom } f$ is defined as

$$\partial f(x) = \{ u \in H \mid (\forall y \in H) \quad f(y) - \langle y - x \mid u \rangle \ge f(x) \}.$$
 (A.6)

and if f is differentiable, $\partial f(x) = \{\nabla f(x)\}$

Theorem A.1 (Fermat's rule). Let $f \in \Gamma_0(\mathcal{H})$. Then:

$$\hat{x} \in \operatorname*{arg\,min}_{x \in \mathcal{H}} f(x) \iff 0 \in \partial f(\hat{x})$$
 (A.7)

Corollary A.2. Let $f, g \in \Gamma_0(\mathcal{H})$ and f is differentiable. Then

$$\hat{x} \in \underset{x \in \mathcal{H}}{\operatorname{arg\,min}} f(x) + g(x) \iff 0 \in \nabla f(\hat{x}) + \partial g(\hat{x})$$
(A.8)

¹Without them, this manuscript would not exist.

These two results are the cornerstone of the convex optimization theory. By studying the variation of f and g we can determine the solution of the optimization problem.

Proposition A.3. Let $f : \mathcal{H} \to \mathbb{R}$ such that $f(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2$ then:

- (i) f is convex and L-smooth.
- (ii) L is the largest singular value of $A^H A$.

(iii)
$$\nabla f(\boldsymbol{x}) = A^H(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y})$$

Definition A.3 (Proximal Operator). Let $f \in \Gamma_0(\mathcal{H})$ and $\lambda > 0$. Then we define the proximal operator (or proximity operator) as

$$\operatorname{prox}_{\lambda f}(\boldsymbol{x}) = \underset{\boldsymbol{u} \in \operatorname{dom} f}{\operatorname{arg\,min}} \frac{1}{2\lambda} \|\boldsymbol{x} - \boldsymbol{u}\|_{2}^{2} + g(\boldsymbol{u})$$
(A.9)

The proximal operator can be considered as the generalization of the projection on convex set operation. Aside from specific cases² there is no closed-form expressions, and one has to use iterative algorithms. Luckily, the proximal operator of usual regularization functions can be expressed in closed form (or are they usual because they have a closed form?).

As we will see later, the regularization function usually takes the form of a scaled norm applied to \boldsymbol{x} , after possibly a change of basis. For complex-valued data, the magnitude is usually considered. The proximal operator of such functions can be expressed using the following proposition:

Proposition A.4 (Complex values). The proximal operator of a function that depends only on the magnitude of its argument can be expressed directly. More precisely, let $f : \mathbb{C} \to \mathbb{R}$, and $g : \mathbb{R} \to \mathbb{R}$ such that, f(z) = g(|z|) then:

$$\operatorname{prox}_{f}(v) = e^{i\angle v} \operatorname{prox}_{q}(|v|) \tag{A.10}$$

Moreover, we recall here the following usual proximal operators:

Proposition A.5 (Usual proximal operators). Let $\lambda > 0$ and \mathcal{H}, \mathcal{G} two Hilbert spaces. The following holds:

(i) Let $\mathcal{H} = \mathbb{R}^N$ and $\mathbf{x} = [x_1, \ldots, x_N] \in \mathcal{H}$, then we can determine the proximal operators for ℓ_2, ℓ_1 and ℓ_0 norms:

$$\operatorname{prox}_{\lambda \|\cdot\|_{2}^{2}}(\boldsymbol{x}) = \frac{1}{1+2\lambda}\boldsymbol{x}$$
(A.11)

$$\operatorname{prox}_{\lambda \parallel \cdot \parallel_1}(\boldsymbol{x}) = \operatorname{soft}_{\lambda}(\boldsymbol{x}) = [\operatorname{sign}(x_i) \max(|x_i| - \lambda, 0)]_{i=1,\dots,N}$$
(A.12)

$$\operatorname{prox}_{\lambda \parallel \cdot \parallel_0}(\boldsymbol{x}) = \operatorname{hard}_{\lambda}(\boldsymbol{x}) = [\operatorname{sign}(x_i) \max(|x_i|, \lambda)]_{i=1,\dots,N}$$
(A.13)

For generalizing to $\mathcal{H} = \mathbb{C}^N$, use $\operatorname{sign}(x_i) = e^{i \angle x_i}$. The soft and hard threshold function are represented in Figure A.1.

²See for instance: https://proximity-operator.net/

(ii) Let $\mathbf{X} \in \mathcal{H} = \mathbb{C}^{N \times M}$ a matrix, and $\mathbf{X} = \sum_{k=1}^{\min(M,N)} s_k \mathbf{u}_k \mathbf{v}_k^H$ its singular values decomposition, then

$$\operatorname{prox}_{\lambda \|\cdot\|_{*}}(\boldsymbol{X}) = \sum_{k=1}^{\min M, N} \operatorname{soft}_{\lambda}(s_{k}) \boldsymbol{u}_{k} \boldsymbol{v}_{k}^{H}$$
(A.14)

(iii) Let $\Psi : \mathcal{H} \to \mathcal{G}$ a linear operator such that $\Psi \Psi^H = \nu \mathrm{Id}_{\mathcal{G}} \ (\nu \neq 0)$, and $g : \mathcal{G} \to \mathbb{R}$, then for $f = g \circ \Psi$:

$$\operatorname{prox}_{\lambda g \circ \Psi}(\boldsymbol{x}) = (\operatorname{Id}_{\mathcal{H}} - \nu^{-1} \Psi^{H} \Psi) \boldsymbol{x} + \nu^{-1} \Psi^{H} \operatorname{prox}_{\lambda \nu g}(\Psi \boldsymbol{x})$$
(A.15)

and if $\mathcal{H} = \mathcal{G}$ (i.e. Ψ is a basis transform) we have:

$$\operatorname{prox}_{\lambda\nu g}(\Psi \boldsymbol{x}) = \nu^{-1} \Psi^{H} \operatorname{prox}_{\lambda\nu g}(\Psi \boldsymbol{x})$$
(A.16)

(iv) (Moreau Identity) Let $g^* = \max_{x \in \mathcal{H}} \langle y, x \rangle + f(x)$ be the Fenchel Legendre conjugate of g^* Then

$$\boldsymbol{x} = \operatorname{prox}_{g}(\boldsymbol{x}) + \operatorname{prox}_{g*}(\boldsymbol{x})$$
 (A.17)

(v) Let $A : \mathcal{H} \to \mathcal{G}$ a linear operator and $f(x) = \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2$. Then

$$\hat{\boldsymbol{u}} = \operatorname{prox}_{f}(\boldsymbol{x}) \iff (\boldsymbol{A}^{*}\boldsymbol{A} + \gamma^{-1}\boldsymbol{I})\hat{\boldsymbol{u}} = (\boldsymbol{A}^{*}\boldsymbol{y} - \boldsymbol{x})$$
 (A.18)

Remark. Eq. (A.16) is typically used with $f = \ell_1$ and Ψ a signal transform such as the (orthogonal) wavelet transform or the Fourier transform. However, it does not hold for redundant wavelet dictionary.

When no closed-form exists, an iterative scheme can be used to



Figure A.1: Hard and soft thresholding functions

A.2 Optimization algorithms

Now that our mathematical toolbox has been stuffed, we can introduce algorithms to solve different cases for the reconstruction.

Mathematically, all the algorithms provide a sequence of elements whose limits are the solution to a well-posed problem. In practice, the algorithm is halted once a sufficient convergence criteria is met, this can be either that the relative change between two iterations is small, or that the cost functions evaluation are reaching a minimum.

A.2.1 Gradient Descent

Consider the following optimization problem:

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}\in\mathcal{H}} f(\boldsymbol{x}) \tag{A.19}$$

Algorithm A.6 (Gradient Descent). Let f a L-smooth function, then one can use the simplest optimization algorithm: gradient descent. It consists of iteratively updating the solution by moving in the opposite direction of the gradient of the function. Let $\eta > 0$. The gradient descent algorithm is:

$$\forall k = 0, \dots N$$

$$\begin{bmatrix} \boldsymbol{x}_{k+1} &= \boldsymbol{x}_k - \eta \nabla f(\boldsymbol{x}_k) \end{bmatrix}$$
(A.20)

Where η is called the step size. This iterative algorithm is guaranteed to converge to a local minimum if the function is convex and $\eta < 1/L$

However, it can be slow to converge, especially in high dimensions. To accelerate the convergence, two methods are commonly used:

Algorithm A.7 (Congugate Gradient). The Conjugate Gradient (CG) method is a variant of the gradient descent algorithm that uses the conjugate gradient of the function to update the solution. In the context where $f(\boldsymbol{x}) = \|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}\|_2^2$, $\boldsymbol{A} \in \mathbb{C}^{N \times M}$ is a data-consistency term (A.4), with $\boldsymbol{p}_0 = \boldsymbol{r}_0 = \boldsymbol{A}^*(\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}_0)$

$$\forall k = 0, \dots N$$

$$\begin{vmatrix} \alpha_k &= \frac{\mathbf{r}_k^H \mathbf{r}_k}{\mathbf{p}_k^H \mathbf{A}^H \mathbf{A} \mathbf{p}_k} \\ \mathbf{r}_{k+1} &= \mathbf{r}_k - \eta_k \mathbf{A}^H \mathbf{A} \mathbf{p}_k \\ \mathbf{p}_{k+1} &= \mathbf{r}_{k+1} + \frac{\mathbf{r}_{k+1}^H \mathbf{r}_{k+1}}{\mathbf{r}_k^H \mathbf{r}_k} \mathbf{p}_k \\ \mathbf{x}_{k+1} &= \mathbf{x}_k - \eta_k \mathbf{p}_k \end{aligned}$$
(A.21)

It can be shown that the Conjugate gradient descent method converges in at most N iterations.

In the absence of any regularization term, the conjugate gradient method is typically used to solve Eq. (A.1), as it is more computationally efficient than trying to use the pseudo-inverse.

A.2.2 Splitting methods

When the objective to minimize is of the form f + g where at least one of the two functions is not differentiable, Gradient Descent methods and their variation are not applicable. To circumvent this, the main approach consists in splitting the problem by reducing the cost of f and g alternatively. In this section, we will introduce the main splitting methods used to solve the optimization problem (A.4). Their usage will be further discussed in the next chapter

Forward-Backward Splitting

Consider now the case where we aim at minimizing f + g where f is L-smooth, and g is proper (but not necessarily differentiable). Since g is not differentiable, we will use the proximal operator of g. This leads to the development of the Forward-Backward Algorithm (presented here in a very generic formulation).

Algorithm A.8 (Forward-Backward, ISTA). Assume that f is β -Lipschitz. Let $\varepsilon \in (0, \min(1, \beta^{-1})), x_0 \in \mathcal{H}$ and $(\lambda_k)_{k \in \mathbb{N}} \in (0, 1)$

 $\forall k = 0, \dots, N \\ \begin{bmatrix} \eta_k \in (\varepsilon, 2/\beta - \varepsilon) \\ \boldsymbol{z}_k = \boldsymbol{x}_k - \eta_k \nabla f(\boldsymbol{x}_k) & // \text{ forward step} \\ \boldsymbol{x}_{k+1} = \boldsymbol{x}_k + \lambda_k (\operatorname{prox}_{\eta_k g}(\boldsymbol{z}_k) - \boldsymbol{x}_k) & // \text{ backward step} \end{bmatrix}$ (A.22)

The sequence $(\boldsymbol{x}_k)_{k \in \mathbb{N}}$ converges to a solution of (A.4).

Remark. If $g = \ell_1$ then the problem is also known as the LASSO, and the Forward Backward algorithm as the Iterative Soft Thresholding Algorithm (ISTA) [BL08; Tib96]. If g = 0, then (A.22) reduces to the gradient descent.

Acceleration

The Forward-Backward Algorithm can be slow to converge, at best it has a linear convergence rate (if the proper assumptions are made on f and g). It is possible to accelerate the convergence of this algorithm by tuning the η_k and embedding some memory in the gradient descent. Notably, the FISTA Algorithm (Fast ISTA)[BT09] uses the following update rule, providing a quadratic convergence:

Algorithm A.9 (FISTA). Let $z_0 = x_0 \in \mathcal{H}$ and use $\eta_k = 1/\beta$

$$\forall k = 0, \dots N$$

$$\begin{aligned} \mathbf{u}_{k} &= \mathbf{z}_{k} - \beta^{-1} \nabla f(\mathbf{z}_{k}) \\ \mathbf{x}_{k+1} &= \operatorname{prox}_{\beta^{-1}g}(u_{k}) \\ t_{k+1} &= \frac{1 + \sqrt{4t_{k}^{2} + 1}}{t_{k+1}} \\ \lambda_{k} &= \frac{t_{k+1} + t_{k} - 1}{t_{k+1}} \\ \mathbf{z}_{k+1} &= \mathbf{x}_{k} + \lambda_{k} (\mathbf{x}_{k+1} - \mathbf{x}_{k}). \end{aligned}$$

$$(A.23)$$

Other acceleration are also possible, and they are closely related to Nesterov acceleration (very similar to FISTA) and POGM, which is adds more memory terms, but converges faster [BT09; KF18]

Backward-Backward Splitting

When both f and g are non-differentiable, we can use the following Backward-Backward splitting:

Algorithm A.10 (Backward-Backward, Half-Quadratic Splitting). Let $z_0 = x_0 \in \mathcal{H}$ and use $\eta_k = 1/\beta$

$$\forall k = 0, \dots N$$

$$\begin{bmatrix} \boldsymbol{u}_k = \operatorname{prox}_f(\boldsymbol{x}_k) \\ \boldsymbol{x}_{k+1} = \operatorname{prox}_g(\boldsymbol{u}_k). \end{bmatrix}$$
(A.24)

Then $(x_k)_{k\in\mathbb{N}}$ does not converge to a solution of (A.4) but to the solution of $f + {}^{\gamma}g$ where ${}^{\gamma}g(x) = \inf_{y\in\mathcal{H}} g(y) + \frac{1}{2} ||x - y||_2^2$ is the Moreau envelope of g.

The Douglas Rachford Splitting extends the Forward-Backward algorithm two to non-differentiable functions.

Algorithm A.11 (Douglas Rachford Splitting). Let $0 < \lambda_k \leq 2$ and $\gamma > 0$. The Douglas-Rachford algorithm is defined as:

$$\forall k = 0, \dots N$$

$$\begin{bmatrix} \boldsymbol{y}_k &= \operatorname{prox}_{\gamma g}(\boldsymbol{x}_k) \\ \boldsymbol{x}_{k+1} &= \boldsymbol{y}_k + \lambda_k (\operatorname{prox}_{\gamma f}(2\boldsymbol{y}_k - \boldsymbol{x}_k) - \boldsymbol{y}_k) \end{bmatrix}$$
(A.25)

The sequence $(\mathbf{x}_k)_{k \in \mathbb{N}}$ converges towards a solution of (A.4).

Primal-Dual algorithm

As we saw in the previous chapter, it is not rare that a sparsity enforcing prior is applied in a specific space \mathcal{G} and g is replaced by $g \circ \Psi$ with $\Psi : \mathcal{H} \to \mathcal{G}$ a linear operator. We now want to solve

$$\hat{\boldsymbol{x}} \arg\min_{\boldsymbol{x}} f(\boldsymbol{x}) + g(\boldsymbol{\Psi}\boldsymbol{x})$$
 (A.26)

Aside from specific case where Theorem A.5 is valid we cannot use any of the presented algorithms.

Algorithm A.12 (Primal Dual Algorithm). Let $\sigma > 0, \tau > 0$ s.t. $\sigma \tau \|\Psi\|^2 \leq 1, \boldsymbol{x}_0 \in \mathcal{H}, \boldsymbol{u}_0 \in \mathcal{G}$

$$\forall k = 0, \dots N$$

$$\begin{bmatrix} \boldsymbol{x}_{k+1} = \operatorname{prox}_{\tau f}(\boldsymbol{x}_k - \tau \Psi^* \boldsymbol{u}_k) \\ \boldsymbol{u}_{k+1} = \operatorname{prox}_{\sigma g^*}(\boldsymbol{u}_k + \sigma \Psi(2\boldsymbol{x}_{k+1} - \boldsymbol{x}_k)) \end{bmatrix}$$
(A.27)

A.2.3 Alternating Direction Method of Multipliers (ADMM)

Another way of solving (A.26) is to rewrite it as

$$\underset{\boldsymbol{u},\boldsymbol{v}\in\mathcal{H}\times\mathcal{G}}{\arg\min} f(\boldsymbol{u}) + g(\boldsymbol{v}), \quad \text{s.t.} \quad \Psi \boldsymbol{u} = \boldsymbol{v},$$
(A.28)

or in its generic form:

$$\underset{\boldsymbol{u},\boldsymbol{v}\in\mathcal{H}\times\mathcal{G}}{\arg\min} f(\boldsymbol{u}) + g(\boldsymbol{v}), \quad \text{s.t.} \quad \boldsymbol{A}\boldsymbol{u} + \boldsymbol{B}\boldsymbol{v} = \boldsymbol{y}, \tag{A.29}$$

where $\mathbf{A}: \mathcal{H} \to \mathcal{J}, \mathbf{B}: \mathcal{G} \to \mathcal{J}$ are two linear operators and $\mathbf{y} \in \mathcal{J}$. Then by setting the augmented Lagrangian, $\tau > 0$ and $\mathbf{\lambda} \in \mathcal{J}$

$$\mathcal{L}_{\tau}(u,v,\lambda) = H(u) + G(v) + \langle \lambda | Au + Bv - b \rangle + \frac{\tau}{2} \| \mathbf{A}u + \mathbf{B}v - b \|^2$$
(A.30)

The ADMM method consists in minimizing L_{τ} alternatively over \boldsymbol{u} and \boldsymbol{v} and the Lagrange multiplier $\boldsymbol{\lambda}$.

Algorithm A.13 (ADMM). Let $u_0 \in \mathcal{H}$, $v_0 \in \mathcal{H}$ and $\tau > 0$

$$\forall k = 0, \dots, N$$

$$\begin{bmatrix} \boldsymbol{u}_{k+1} = \arg\min_{\boldsymbol{u}} \mathcal{L}_{\tau}(\boldsymbol{u}, \boldsymbol{v}_{k}, \boldsymbol{\lambda}_{k}) \\ \boldsymbol{v}_{k+1} = \arg\min_{\boldsymbol{v}} \mathcal{L}_{\tau}(\boldsymbol{u}_{k+1}, \boldsymbol{v}, \boldsymbol{\lambda}_{k}) \\ \lambda_{k+1} = \boldsymbol{\lambda}_{k} + \tau(\boldsymbol{A}\boldsymbol{u} + \boldsymbol{B}\boldsymbol{v} - \boldsymbol{b}) \end{bmatrix}$$
(A.31)

Moreover, by using a re-scaled Lagrange multiplier $\boldsymbol{\mu} = \tau^{-1} \boldsymbol{\lambda}$ we can rewrite the two first step of ADMM in a f + g form

$$\forall k = 0, \dots, N$$

$$\begin{aligned} \mathbf{u}_{k+1} &= \arg\min_{\boldsymbol{u}} \tilde{\mathcal{L}}_{\tau}(\boldsymbol{u}, \boldsymbol{v}_{k}, \boldsymbol{\mu}_{k}) \\ &= \arg\min_{\boldsymbol{u}} f(\boldsymbol{u}) + \frac{\tau}{2} \|\boldsymbol{A}\boldsymbol{u} + \boldsymbol{B}\boldsymbol{v}_{k} - \boldsymbol{b} + \boldsymbol{\mu}_{k}\|_{2}^{2} \\ \boldsymbol{v}_{k+1} &= \arg\min_{\boldsymbol{v}} \tilde{\mathcal{L}}_{\tau}(\boldsymbol{u}_{k+1}, \boldsymbol{v}, \boldsymbol{\mu}_{k}) \\ &= \arg\min_{\boldsymbol{v}} g(\boldsymbol{v}) + \frac{\tau}{2} \|\boldsymbol{A}\boldsymbol{u}_{k+1} + \boldsymbol{B}\boldsymbol{v} - \boldsymbol{b} + \boldsymbol{\mu}_{k}\|_{2}^{2} \\ \boldsymbol{\mu}_{k+1} &= \boldsymbol{\mu}_{k} + (\boldsymbol{A}\boldsymbol{u}_{k+1} + \boldsymbol{B}\boldsymbol{v}_{k+1} - \boldsymbol{y}) \end{aligned}$$
(A.32)

the sequence (u_k and v_k will converge to a saddle point of the Lagrangian L, and more precisely will minimize f + g.

Remark. ADMM is a very generic formulation, and requires solving two internal optimization problems at each iteration. It can also be noted that by setting $\mathbf{A} = -\mathbf{B} = Id_H$ and $\mathbf{b} = 0$ we get the Forward-algorithm back (Theorem A.8)

Acceleration

ADMM is a notoriously slow algorithm. Indeed, each step requires solving internal optimization problems, that do not necessarily have closed form. To overcome this, an accelerated scheme for ADMM have been proposed [Gol+14], using the same techniques developed for accelerating the Forward-Backward algorithm. The internal optimization steps are typically solved with theorem A.8, which can also be accelerated. Finally, a practical speed up can be obtained by *warm-starting* the inner optimization by their previous results. Indeed, it can be assumed that the last iteration is a reasonable first guess, and will position the start of the algorithm close to the look-after optimal point, reducing the number of steps required in practice.

A.2.4 More than two functions

So far, we have restricted the problem to be the sum of two proper convex functions f + g, but what if the problem at hand is further regularized by using multiple priors such that $g = g_1 + g_2 + \cdots + g_R$. However, the proximal operator of the sum of R function (even when R = 2) is not easy to compute.

Luckily enough, mathematicians have proposed solutions for this problem as well [PCP11; CP21]. In practice, the regularization terms consist in only 2 to 3 components, and algorithms are designed specifically in these cases.

Appendix B

Appendix for Chapter 5

B.1 BrainWeb Phantom



(a) Fuzzy tissues masks at 0.rmm iso resolution.



(b) TE=10ms, TR=50ms, $FA=45^{\circ}$.



(c) TE=30ms, TR=50ms, $FA=12^{\circ}$.

(d) Tissue Parameters used for the simulation, relaxations times (T_1, T_2, T_2^*) are given in ms at 7T), proton density (ρ) and magnetic susceptibility χ are dimensionless.

Tissue	T_1	T_2	T_2^*	ρ	χ
WM	1200	57	27	0.77	-9.08
GM	1800	49	28	0.86	-9.05
CSF	3730	1010	1010	1	-9.05

Figure B.1: BrainWeb phantom [Aub+06] and tissue parameters used for the simulations.

B.2 Estimation of the regularization parameter using SURE.

In § 7.5.2 we explained that the regularization parameter μ_t used in Eq. (7.10) was estimated for each frame t. In practice, this is achieved as described in ?? 1. As the estimate x_t becomes cleaner, the threshold automatically decreases. This estimation is very fast to compute because we restrict the estimation to the most detailed wavelet coefficients in the high-frequency subbands. Therefore, the problem dimension is reduced and we can use a line search strategy to find $\hat{\mu}_t$ without noticeable overhead for image reconstruction.

 $\begin{array}{l} \text{Data: } x_t, \Psi\\ \text{Result: } \mu_t\\ \alpha_H \leftarrow P^{HHH}(\Psi(x_t)) & // \text{ Extract high-details wavelet coeffs} \\ n \leftarrow \text{size}(\alpha_H)\\ \sigma \leftarrow \text{MAD}(\alpha_H) * 0.675 & // \text{ Estimate variance using Median Absolute}\\ \text{Deviation} \\ \alpha_H \leftarrow \alpha_H / \sigma\\ \text{if } \frac{1}{n} \alpha_H^T \alpha_H < \log_2(n)^{3/2} / \sqrt{n} \text{ then} \\ \mid \mu_t \leftarrow \sqrt{2 \log_2(n)} \\ \text{else} \\ \mid \hat{\mu}_t \leftarrow \arg\min_{w \in \alpha_H} \sum_{i=1}^n \min(\alpha_{H,i}^2, w^2) - 2 * I\{\alpha_{H,i}^2 < w^2\} & // \text{ Find a threshold} \\ \text{ that minimizes the SURE estimator} \\ \mid \mu_t \leftarrow \min(\hat{\mu}_t, \sqrt{2 \log_2(n)}) \\ \text{end} \end{array}$

Algorithm 1: Estimation of μ_t for solving (7.10)

B.3 Temporal SNR maps for the three scenarios

To ensure that the values for SNR_i have been correctly chosen, here we report the tSNR maps for each scenario. The tSNR in absence of evoked brain activity is defined voxel-wise as:

$$\forall v \in \text{FOV}, \quad tSNR(v) = \text{mean}(x_t(v))/\operatorname{std}(x_t(v)), \tag{B.1}$$

where $x_t(v)$ refers to the BOLD fMRI time series at frame t and location v.

Interestingly, in Figure B.2a we show that the tSNR is marginally impacted by T_2^* relaxation effects in Scenario 1 due to the use of Cartesian 3D EPI. In regards to Scenario 2 we already reported that T_2^* relaxation effects were detrimental to the mean tSNR, mainly due to longer T_{obs} . This is confirmed in Figure B.2b with lower tSNR values compared

to those shown in Figure B.2a while both scenarios were initialized with the same input SNR and defined at the same spatial resolution. Furthermore, one can observe that the static refined strategy maintains a better tSNR than the dynamic refined one, but is associated instead with a more degraded image quality (significant blurring). Figure B.2c depicts the tSNR map for Scenario 3 under the T_2^* model, which lies in the same range as for S2 while being smoother spatially.



(a) tSNR maps for scenario S1. Top: T_2^* model; Bottom: simple Fourier model. 20 40 60 ^{5NR} 80

(c) tSNR map for scenario S3 with T_2^* model.



Figure B.2: tSNR maps for the 3 scenarios.

List of Tables and Figures

1	Software developed during this thesis
$\begin{array}{c} 1.1 \\ 1.2 \end{array}$	The neurovascular coupling and the hemodynamic response function Relaxation of the magnetization vector and the Free Induction Decay (FID) signal
1.3	EPI2D and Spiral in k-space and their gradient profiles
1.4	BOLD sensitivity and its impact on TE
1.5	Pulse sequence diagram for a GRE 3D EPI.
1.0	MOCK IMRI experiment.
1.1	BOLD time series example
1.0	Example of activation maps for auditory stimulus
1.10	fMRI processing chain pyramid.
2.1	Multi-coil array description.
2.2	3D Non-Cartesian undersampling patterns for MRI acquisition.
2.3	TSNR map comparing EPI and spiral based trajectory, from [Glo12]. The spiral trajectory shows a better tSNR in the frontal and temporal lobes
2.4	TURBINE vs EPI
2.5	Example of compressed sensing reconstruction in the 2D Cartesian case.
2.6	CS MRI components
4.1	U-Net Architecture. Adapted from [Zha+21b]
4.2	Neural Network for Inverse problems
5.1	MRI-NUFFT as an interface for non-Cartesian MRI
5.2	Equivalence between NUFFT operations
5.3	Density compensation for radial trajectory.
5.4	Benchmark of NUFFT
5.5	Effect of the field inhomegeneities correction on the fMRI data
6.1	General procedure for LLR denoising
6.2	LLR methods parametrization.
6.3	Non-zero Z-score probability distribution after denoising
6.4	Z-score activation maps
6.5	Zoom Activations maps for different denoising methods
6.6	Average gain (multiplicative factor) \overline{MF} of activated voxels using denoising methods compared to baseline.
71	Summary of characteristics of different fMRI data simulators

List of Tables and Figures

7.2	Acquisition method implemented in SNAKE – The case represented is simplified to a 2D Cartesian case (<i>e.g.</i> a projected view of a 3D non- accelerated EPI scheme). Each shot (i.e. a plane in 3D EPI) of the k-space sampling pattern is acquired separately from an on-the-fly simulated volume in the image domain as shown in the blue frame. The shots are numbered here from ① to ⑦. The parallel acquisition is performed in parallel for each tissue type to apply the T_2^* relaxation model (7.5)	99
7.3	handlers-chaining	101
7.4	Overview of simulated scenarios and their computation requirements – L, N_s, n_{jobs} specify the number of simulated coils, the number of shots acquired to create a k-space volume and the number of concurrently simulated shots, respectively.	104
7.5	Non-Cartesian k-space sampling trajectories used in Scenarios 2 and 3.	105
7.6	Different methodologies for sequential reconstruction used in Scenario S2. Top: <i>cold start</i> reconstruction, each frame is reconstructed independently. Center: <i>warm start</i> reconstruction, each frame is reconstructed using the previous frame as initialization. Bottom: <i>refined</i> reconstruction, after a warm start reconstruction, the last frame is used as initialization for all other frames.	108
7.7	Comparison of the two acquisition models (basic Fourier and T_2^* effects) for scenario S1.	109
7.8	Activations maps for scenario S2. Top: T_2^* relaxation is taken into account. Bottom: Basic Fourier model, no T_2^* relaxation. Colored frames within each insert follow the convention adopted in the tables reported in Figure 7.9. Detected activations surviving at $p < 0.001$, uncorrected, thresholding are	110
7.9	Quantitative metrics summarizing image quality (SSIM/PSNR), signal quality (tSNR) and statistical performances (AUC/BACC) for scenario S2. Top: T_2^* relaxation is taken into account. Bottom: Basic Fourier model. TSNR values are averaged over the occipital ROI. In the tables, First and Last refer to the first and last frames in the fMRI volumes sequences	110
7.10	Results for scenario S3. Top left/center: First fMRI reconstructed volume using the standard Fourier (left) and T_2^* (center) models. Sequential cold- start CS reconstruction was performed frame-wise. Detected activations surviving at $p < 0.001$ thresholding are overlaid using a colorjet map. Top right: Precision/Recall curves for the two models. Accurate modeling of realistic T_2^* decay slightly impacts statistical performances. Bottom: Quantitative metrics summarizing image quality (SSIM/PSNR), signal quality (tSNR) and statistical performances (AUC/BACC)	111
		110
8.1	Effect of the unsupervised preprocessing on the training dataset	125
8.2	Benchmark PnP Results	127
8.3	Reconstruction metrics for different solvers and acceleration factors	128
8.4	Evolution of the PSNR across iterations for various PnP algorithms at AF=16.	129

9.1	Brain Phantom used in our simulation (top) and average BOLD re- sponse (simulated and measured) in the ROI (bottom).	133
9.2	Stack of spirals used to sample a single frame in 3D k-space. This pattern is repeated over time to collect the full fMBI sequence	135
9.3	Z-score maps for CG/CS/HQS-F1 reconstruction methods (from left to right). the target ROI (voxel with at least 50% of gray matter, and belonging to primary visual cortex) is delineated in cvan. The complete	100
0.4	anatomical ROI (occipital cortex) is contoured in green.	136
9.4	Quantitative metrics on image quality and statistical performances. SSIM and PSNR scores are averaged over all frames. The balanced accuracy (BACC) is computed with a threshold set to $p < 10^{-3}$ and AUC and BACC	
9.5	were computed with respect to voxels with at least 50% of gray matter. Precision-Recall characteristic curve for the different thresholds of gray	137
	matter considered in the ROI	137
A.1	Hard and soft thresholding functions	151
B.1 B.2	BrainWeb phantom [Aub+06] and tissue parameters used for the simulations tSNR maps for the 3 scenarios.	$\begin{array}{c} 157 \\ 160 \end{array}$

Bibliography

- [Abr+14] Alexandre Abraham et al. "Machine Learning for Neuroimaging with Scikit-Learn". In: *Frontiers in Neuroinformatics* 8 (2014). ISSN: 1662-5196 (pages 84, 101, 102).
- [AO18] Jonas Adler and Ozan Oktem. "Learned Primal-dual Reconstruction". In: *IEEE Transactions on Medical Imaging* 37.6 (June 2018), pp. 1322–1332. ISSN: 0278-0062, 1558-254X (pages 60, 62).
- [AMJ19] Hemant Kumar Aggarwal, Merry P. Mani, and Mathews Jacob. "MoDL: Model Based Deep Learning Architecture for Inverse Problems". In: *IEEE Transactions* on Medical Imaging 38.2 (Feb. 2019), pp. 394–405. ISSN: 0278-0062, 1558-254X (pages 58, 60).
- [Ahm+20] Rizwan Ahmad et al. "Plug and Play Methods for Magnetic Resonance Imaging (Long Version)". In: *IEEE Signal Processing Magazine* 37.1 (Jan. 2020), pp. 105– 116. ISSN: 1053-5888, 1558-0792 (page 122).
- [Akç+19] Mehmet Akçakaya et al. "Scan-Specific Robust Artificial-Neural-Networks for k-Space Interpolation (RAKI) Reconstruction: Database-free Deep Learning for Fast Imaging". In: *Magnetic Resonance in Medicine* 81.1 (Jan. 2019), pp. 439– 453. ISSN: 1522-2594 (page 36).
- [Akç+21] Mehmet Akçakaya et al. "Unsupervised Deep Learning Methods for Biological Image Reconstruction". May 17, 2021 (pages 59, 64).
- [Amo+22a] Z. Amor et al. "Prospects of Non-Cartesian 3D-SPARKLING Encoding for Functional MRI: A Preliminary Case: Study for Retinotopic Mapping". In: Proceedings of the 30'th Scientific Meeting of the International Society for Magnetic Resonance in Medicine (2022) (pages 80, 114).
- [Amo+23] Z. Amor et al. "Non-Cartesian Non-Fourier FMRI Imaging for High-Resolution Retinotopic Mapping at 7 Tesla". In: 2023 IEEE 9th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP). Dec. 2023, pp. 201–205 (pages 9, 30, 43, 45, 69, 76, 77, 97, 104).
- [Amo24] Zaineb Amor. "Non-Cartesian Sparkling Encoding for High Spatio-Temporal Resolution Functional Magnetic Resonance Imaging (fMRI) at 7 Tesla and Beyond". These de doctorat. université Paris-Saclay, Apr. 9, 2024 (pages 16, 36, 45).
- [Amo+22b] Zaineb Amor et al. "B0 Field Distortions Monitoring and Correction for 3D Non-Cartesian fMRI Acquisitions Using a Field Camera: Application to 3D-SPARKLING at 7T". In: Joint Annual Meeting ISMRM-ESMRMB & ISMRT 31st Annual Meeting. May 7, 2022 (page 43).
- [Amo+24a] Zaineb Amor et al. "Impact of B0\ \mathrmB_0 \ Field Imperfections Correction on BOLD Sensitivity in 3D-SPARKLING fMRI Data". In: *Magnetic Resonance in Medicine* 91.4 (2024), pp. 1434–1448. ISSN: 1522-2594 (pages 30, 93, 104).

Bibliography

[Amo+24b]	Zaineb Amor et al. "Non-Cartesian 3D-SPARKLING vs Cartesian 3D-EPI Encoding Schemes for Functional Magnetic Resonance Imaging at 7 Tesla". In: <i>PLOS ONE</i> 19.5 (May 13, 2024), e0299925 (page 45).
[Amo+24c]	Zaineb Amor et al. "Non-cartesian 3D-SPARKLING vs cartesian 3D-EPI encod- ing schemes for functional magnetic resonance Imaging at 7 Tesla". In: <i>Plos one</i> 19.5 (2024), e0299925 (pages 105, 106, 113, 115).
[Ash19]	F. Gregory Ashby. <i>Statistical Analysis of fMRI Data, Second Edition.</i> MIT Press, Sept. 17, 2019. 569 pp. ISBN: 978-0-262-04268-0 (pages 16, 32).
[AE91]	D. J. Atkinson and R. R. Edelman. "Cineangiography of the Heart in a Single Breath Hold with a Segmented turboFLASH Sequence". In: <i>Radiology</i> 178.2 (Feb. 1991), pp. 357–360. ISSN: 0033-8419 (page 45).
[Aub+06]	B. Aubert-Broche et al. "Twenty New Digital Brain Phantoms for Creation of Validation Image Data Bases". In: <i>IEEE Transactions on Medical Imaging</i> 25.11 (Nov. 2006), pp. 1410–1416. ISSN: 1558-254X (page 157).
[BVC13]	Solveig Badillo, Thomas Vincent, and Philippe Ciuciu. "Group-Level Impacts of within- and between-Subject Hemodynamic Variability in fMRI". In: <i>NeuroImage</i> 82 (Nov. 15, 2013), pp. 433–448. ISSN: 1053-8119 (page 17).
[Ban+92]	P. A. Bandettini et al. "Time Course EPI of Human Brain Function during Task Activation". In: <i>Magnetic Resonance in Medicine</i> 25.2 (June 1992), pp. 390–397. ISSN: 0740-3194 (pages 1, 5).
[BHF21]	Peter A Bandettini, Laurentius Huber, and Emily S Finn. "Challenges and opportunities of mesoscopic brain mapping with fMRI". In: <i>Current Opinion in Behavioral Sciences</i> 40 (2021), pp. 189–200 (page 132).
[BMK19]	Alex H. Barnett, Jeremy F. Magland, and Ludvig af Klinteberg. "A Parallel Non- Uniform Fast Fourier Transform Library Based on an "Exponential of Semicircle" Kernel". Apr. 8, 2019 (page 70).
[BP11]	Markus Barth and Benedikt A. Poser. "Advances in High-Field BOLD fMRI". In: <i>Materials</i> 4.11 (Nov. 2, 2011), pp. 1941–1955. ISSN: 1996-1944 (page 80).
[Bar+15]	Markus Barth et al. "Simultaneous Multislice (SMS) Imaging Techniques". In: Magnetic Resonance in Medicine 75.1 (Aug. 26, 2015), p. 63 (pages 26, 37).
[Bau+17]	Heinz H Bauschke et al. Correction to: convex analysis and monotone operator theory in Hilbert spaces. 2017 (page 126).
[BC11]	Heinz H. Bauschke and Patrick L. Combettes. <i>Convex Analysis and Monotone Operator Theory in Hilbert Spaces.</i> CMS Books in Mathematics. New York, NY: Springer New York, 2011. ISBN: 978-1-4419-9466-0 978-1-4419-9467-7 (page 149).
[BT09]	Amir Beck and Marc Teboulle. "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems". In: <i>SIAM Journal on Imaging Sciences</i> 2.1 (Jan. 2009), pp. 183–202. ISSN: 1936-4954 (pages 106, 153).
[BMW09]	CM Bennett, MB Miller, and GL Wolford. "Neural Correlates of Interspecies Perspective Taking in the Post-Mortem Atlantic Salmon: An Argument for Multiple Comparisons Correction". In: <i>NeuroImage</i> . Organization for Human Brain Mapping 2009 Annual Meeting 47 (July 1, 2009), S125. ISSN: 1053-8119 (page 32).

[BM10]	Craig M. Bennett and Michael B. Miller. "How Reliable Are the Results from Functional Magnetic Resonance Imaging?" In: <i>Annals of the New York Academy of Sciences</i> 1191 (Mar. 2010), pp. 133–155. ISSN: 1749-6632 (page 93).
[Ber+14]	Michaël Bernier et al. "Using fMRI Non-Local Means Denoising to Uncover Activation in Sub-Cortical Structures at 1.5 T for Guided HARDI Tractography". In: <i>Frontiers in Human Neuroscience</i> 8 (Sept. 11, 2014), p. 715. ISSN: 1662-5161 (page 83).
[Bhu+16]	Chitresh Bhushan et al. "Temporal Non-Local Means Filtering Reveals Real-Time Whole-Brain Cortical Interactions in Resting fMRI". In: <i>PLoS ONE</i> 11.7 (July 8, 2016), e_504_0158. ISSN: 1932-6203 (page 83).
[Blo46]	F. Bloch. "Nuclear Induction". In: Physical Review 70.7-8 (Oct. 1, 1946), pp. 460–474 (page 96).
[Bol+17]	Saskia Bollmann et al. "Analysis and Correction of Field Fluctuations in fMRI Data Using Field Monitoring". In: <i>NeuroImage</i> . Cleaning up the fMRI Time Series: Mitigating Noise with Advanced Acquisition and Correction Strategies 154 (July 1, 2017), pp. 92–105. ISSN: 1053-8119 (pages 72, 93).
[Bos+20]	Han Bossier et al. "The empirical replicability of task-based fMRI as a function of sample size". In: <i>NeuroImage</i> 212 (2020), p. 116601 (page 93).
[BV04]	Stephen P. Boyd and Lieven Vandenberghe. <i>Convex Optimization</i> . Cambridge, UK ; New York: Cambridge University Press, 2004. 716 pp. ISBN: 978-0-521-83378-3 (page 149).
[BL08]	Kristian Bredies and Dirk A. Lorenz. "Linear Convergence of Iterative Soft-Thresholding". In: <i>Journal of Fourier Analysis and Applications</i> 14.5-6 (Dec. 2008), pp. 813–837. ISSN: 1069-5869, 1531-5851 (page 153).
[Bre+05]	Felix A. Breuer et al. "Controlled Aliasing in Parallel Imaging Results in Higher Acceleration (CAIPIRINHA) for Multi-Slice Imaging". In: <i>Magnetic Resonance in Medicine</i> 53.3 (2005), pp. 684–691. ISSN: 1522-2594 (pages 37, 45, 115).
[Bre+09]	Felix A. Breuer et al. "General Formulation for Quantitative G-factor Calculation in GRAPPA Reconstructions". In: <i>Magnetic Resonance in Medicine</i> 62.3 (Sept. 2009), pp. 739–746. ISSN: 07403194, 15222594 (page 81).
[Bro+14]	Robert W. Brown et al. <i>Magnetic Resonance Imaging: Physical Principles and Sequence Design.</i> 1st ed. Wiley, Apr. 22, 2014. ISBN: 978-0-471-72085-0 978-1-118-63395-3 (pages 16, 18, 21).
[BF97]	R.B. Buxton and L.R. Frank. "A Model for the Coupling between Cerebral Blood Flow and Oxygen Metabolism during Neural Stimulation". In: <i>Journal of Cerebral Blood Flow and Metabolism</i> 17.1 (1997), pp. 64–72 (page 16).
[Bux+04]	Richard B. Buxton et al. "Modeling the Hemodynamic Response to Brain Activation". In: <i>NeuroImage</i> 23 (2004), S220–S233. ISSN: 10538119 (page 24).
[Byc+18]	Clare Bycroft et al. "The UK Biobank Resource with Deep Phenotyping and Genomic Data". In: <i>Nature</i> 562.7726 (Oct. 2018), pp. 203–209. ISSN: 1476-4687 (page 6).
[CR17]	César Caballero-Gaudes and Richard C. Reynolds. "Methods for Cleaning the BOLD fMRI Signal". In: <i>NeuroImage</i> . Cleaning up the fMRI Time Series: Mitigating Noise with Advanced Acquisition and Correction Strategies 154 (July 1, 2017), pp. 128–149. ISSN: 1053-8119 (page 93).

Bibliography

[CJO23]	Caglar Cakan, Nikola Jajcay, and Klaus Obermayer. "Neurolib: A Simulation Framework for Whole-Brain Neural Mass Modeling". In: <i>Cognitive Computation</i> 15.4 (July 1, 2023), pp. 1132–1152. ISSN: 1866-9964 (page 95).
[CA12]	V.D. Calhoun and T. Adali. "Analysis of Complex-Valued Functional Magnetic Resonance Imaging Data: Are We Just Going through a "Phase"?" In: <i>Bulletin of the Polish Academy of Sciences: Technical Sciences</i> 60.3 (Dec. 1, 2012), pp. 371–418. ISSN: 0239-7528 (page 21).
[CV23]	Virginie Callot and Alexandre Vignaud. "Imagerie à Haut Champ". In: Hélène Ratiney and Olivier Beuf. <i>Les enjeux de l'IRM</i> . ISTE Edition. 2023. ISBN: 978-1-78948-113-6 (page 24).
[CRT06]	E.J. Candes, J. Romberg, and T. Tao. "Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information". In: <i>IEEE Transactions on Information Theory</i> 52.2 (Feb. 2006), pp. 489–509. ISSN: 1557-9654 (page 40).
[CST13]	Emmanuel J. Candès, Carlos A. Sing-Long, and Joshua D. Trzasko. "Unbiased Risk Estimates for Singular Value Thresholding and Spectral Estimators". In: <i>IEEE Transactions on Signal Processing</i> 61.19 (Oct. 2013), pp. 4643–4657. ISSN: 1941-0476 (page 81).
[Cas+18]	 B. J. Casey et al. "The Adolescent Brain Cognitive Development (ABCD) Study: Imaging Acquisition across 21 Sites". In: <i>Developmental Cognitive Neuroscience</i>. The Adolescent Brain Cognitive Development (ABCD) Consortium: Rationale, Aims, and Assessment Strategy 32 (Aug. 1, 2018), pp. 43–54. ISSN: 1878-9293 (pages 2, 6).
[Cas+23]	Carlos Castillo-Passi et al. "KomaMRI.Jl: An Open-Source Framework for General MRI Simulations with GPU Acceleration". In: <i>Magnetic Resonance in</i> <i>Medicine</i> 90.1 (2023), pp. 329–342. ISSN: 1522-2594 (pages 93–95).
[CRC21]	G R Chaithya, Zaccharie Ramzi, and Philippe Ciuciu. "Learning the Sampling Density in 2D SPARKLING MRI Acquisition for Optimized Image Reconstruction". In: 2021 29th European Signal Processing Conference (EUSIPCO). 2021 29th European Signal Processing Conference (EUSIPCO). Aug. 2021, pp. 960–964 (page 43).
[Cha+22]	G. R. Chaithya et al. "Optimizing Full 3D SPARKLING Trajectories for High-Resolution Magnetic Resonance Imaging". In: <i>IEEE Transactions on Medical Imaging</i> 41.8 (Aug. 2022), pp. 2105–2117. ISSN: 1558-254X (pages 38, 43, 72, 74, 105, 122).
[CWE16]	Stanley H. Chan, Xiran Wang, and Omar A. Elgendy. <i>Plug-and-Play ADMM</i> for Image Restoration: Fixed Point Convergence and Applications. Nov. 11, 2016. URL: http://arxiv.org/abs/1605.01710 (visited on 12/12/2024). Pre-published (page 62).
[CJ23]	Jyothi Rikhab Chand and Mathews Jacob. <i>Multi-Scale Energy (MuSE) Plug and Play Framework for Inverse Problems</i> . May 8, 2023. URL: http://arxiv.org/abs/2305.04775. Pre-published (page 124).
[Cha+14]	Nicolas Chauffert et al. "Variable Density Sampling with Continuous Trajectories". In: <i>SIAM Journal on Imaging Sciences</i> 7.4 (Jan. 2014), pp. 1962–1992 (page 42).
[Che+06]	Hu Cheng et al. "SmartPhantom – an fMRI Simulator". In: <i>Magnetic Resonance Imaging</i> 24.3 (Apr. 1, 2006), pp. 301–313. ISSN: 0730-725X (page 93).
-----------	--
[Che+18]	H. Cherkaoui et al. "Analysis vs Synthesis-based Regularization for Combined Compressed Sensing and Parallel MRI Reconstruction at 7 Tesla". In: 2018 26th European Signal Processing Conference (EUSIPCO) (Sept. 2018), pp. 36–40 (page 44).
[CGM18]	Mark Chiew, Nadine N. Graedel, and Karla L. Miller. "Recovering Task fMRI Signals from Highly Under-Sampled Data with Low-Rank and Temporal Subspace Constraints". In: <i>NeuroImage</i> 174 (July 1, 2018), pp. 97–110. ISSN: 1053-8119 (pages 52, 53, 105).
[Chi+15]	Mark Chiew et al. "K-t FASTER: Acceleration of Functional MRI Data Acquisition Using Low Rank Constraints". In: <i>Magnetic Resonance in Medicine</i> 74.2 (Aug. 2015), pp. 353–364. ISSN: 1522-2594 (pages 39, 45, 51).
[Chi+16]	Mark Chiew et al. "Accelerating Functional MRI Using Fixed-rank Approximations and Radial-cartesian Sampling". In: <i>Magnetic Resonance in Medicine</i> 76.6 (Dec. 2016), pp. 1825–1836. ISSN: 0740-3194, 1522-2594 (page 51).
[CLY22]	Hyungjin Chung, Eun Sun Lee, and Jong Chul Ye. "MR image denoising and super-resolution using regularized reverse diffusion". In: <i>IEEE Transactions on Medical Imaging</i> 42.4 (2022), pp. 922–934 (page 124).
[Ciu+03]	P. Ciuciu et al. "Unsupervised Robust Nonparametric Estimation of the Hemo- dynamic Response Function for Any Fmri Experiment". In: <i>IEEE Transactions</i> on Medical Imaging 22.10 (Oct. 2003), pp. 1235–1251. ISSN: 0278-0062 (page 17).
[CP11a]	Patrick L Combettes and Jean-Christophe Pesquet. "Proximal splitting methods in signal processing". In: <i>Fixed-point algorithms for inverse problems in science and engineering</i> (2011), pp. 185–212 (pages 123, 136, 149).
[CP11b]	Patrick L Combettes and Jean-Christophe Pesquet. "Proximal splitting methods in signal processing". In: <i>Fixed-point algorithms for inverse problems in science and engineering</i> (2011), pp. 185–212 (page 134).
[CP21]	Patrick L. Combettes and Jean-Christophe Pesquet. "Fixed Point Strategies in Data Science". Jan. 25, 2021 (pages 149, 156).
[CVC24a]	PA. Comby , A. Vignaud, and P. Ciuciu. <i>SNAKE-fMRI: A Modular fMRI Data Simulator from the Space-Time Domain to k-Space and Back.</i> Version 2. Apr. 12, 2024. URL: http://arxiv.org/abs/2404.08282 (visited on 09/01/2024). In revision to Imaging Neuroscience (pages 3, 7, 9, 92).
[CVC24b]	PA. Comby , A. Vignaud, and P. Ciuciu. "SNAKE-fMRI: A Modular fMRI Simulator from the Space-Time Domain to k-Space Data and Back". In: <i>ISMRM Annual Meeting, (in Press).</i> Singapore, 2024. Awarded the 2nd best abstract award from the reproducibility study group. (Pages 9, 92, 132, 134).
[Com+23a]	PA. Comby et al. "Benchmarking Local Low Rank Denoising Methods for Task-Based fMRI Data Analysis". In: <i>ISMRM 2023 Annual Meeting</i> . Toronto, June 2023 (page 9).
[Com+23b]	PA. Comby et al. "Denoising of FMRI Volumes Using Local Low Rank Methods". In: ISBI 2023 - International Symposium on Biomedical Imaging. Apr. 18, 2023 (pages 3, 7, 9, 79).

- [Com+24a] P.-A. Comby et al. "MRI-NUFFT: An Open Source Python Package to Make Non-Cartesian MR Imaging Easier". In: ISMRM Annual Meeting, (in Press). Singapore, 2024 (pages 9, 69).
- [Com+24b] P.-A. Comby et al. Robust Plug-and-Play Methods for Highly Accelerated Non-Cartesian MRI Reconstruction. Version 2. Oct. 2024. URL: https://hal. science/hal-04759015. Accepted to ISBI 2025 (pages 9, 121, 135).
- [Com+24c] Pierre-Antoine Comby et al. MRI-NUFFT: Doing Non-Cartesian MRI Has Never Been Easier. Version 2. Nov. 2024. URL: https://hal.science/hal-04775262 (visited on 12/12/2024). in revision to Journal of Open Source Software. (Pages 9, 69).
- [Com+25a] **Pierre-Antoine Comby** et al. *Plug-and-Play Reconstruction for 3D Non-Cartesian fMRI Data.* Mar. 2025. URL: https://hal.science/hal-04993532. Submitted to EUSIPCO 2025 (pages 9, 131).
- [Com+25b] **Pierre-Antoine Comby** et al. "Robust plug-and-play methods for highly accelerated non-Cartesian MRI reconstruction". In: *ISBI 2025* (2025) (page 134).
- [Cox96] R. W. Cox. "AFNI: Software for Analysis and Visualization of Functional Magnetic Resonance Neuroimages". In: Computers and Biomedical Research, an International Journal 29.3 (June 1996), pp. 162–173. ISSN: 0010-4809 (page 29).
- [Dab+22] Arwa Dabbech et al. "First AI for deep super-resolution wide-field imaging in radio astronomy: unveiling structure in ESO 137-006". In: *The Astrophysical Journal Letters* 939.1 (2022), p. L4 (page 123).
- [Dav22a] Guillaume Daval-Frerot. "Deep Neural Networks for MR Image Reconstruction and B0 Inhomogeneity Correction in Non-Cartesian Susceptibility Weighted Imaging at 3 Tesla". These de doctorat. université Paris-Saclay, Dec. 16, 2022 (pages 18, 25).
- [Dav22b] Guillaume Daval-Frerot. "Réseaux de neurones profonds pour la reconstruction en IRM et la correction d'inhomogénéités de B0 de l'imagerie pondérée en susceptibilité magnétique non cartésienne à 3 Tesla". PhD thesis. Université Paris-Saclay, Dec. 16, 2022 (page 16).
- [Dav+22] Guillaume Daval-Frérot et al. "Iterative Static Field Map Estimation for Off-Resonance Correction in Non-Cartesian Susceptibility Weighted Imaging". In: *Magnetic Resonance in Medicine* 88.4 (2022), pp. 1592–1607. ISSN: 1522-2594 (page 72).
- [Des+12] Anagha Deshmane et al. "Parallel MR Imaging". In: Journal of Magnetic Resonance Imaging 36.1 (2012), pp. 55–72. ISSN: 1522-2586 (page 36).
- [DJ94] D.L. Donoho and I.M. Johnstone. "Threshold Selection for Wavelet Shrinkage of Noisy Data". In: Proceedings of 16th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. 16th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. Baltimore, MD, USA: IEEE, 1994, A24–A25. ISBN: 978-0-7803-2050-5 (page 106).
- [Dro+06] Ivana Drobnjak et al. "Development of a Functional Magnetic Resonance Imaging Simulator for Modeling Realistic Rigid-Body Motion Artifacts". In: *Magnetic Resonance in Medicine* 56.2 (2006), pp. 364–380. ISSN: 1522-2594 (pages 95, 96, 105).

- [Due+15] Yolanda Duerst et al. "Real-Time Feedback for Spatiotemporal Field Stabilization in MR Systems". In: *Magnetic Resonance in Medicine* 73.2 (2015), pp. 884–893. ISSN: 1522-2594 (page 93).
- [Duy12] Jeff H. Duyn. "The Future of Ultra-High Field MRI and fMRI for Study of the Human Brain". In: *NeuroImage*. 20 YEARS OF fMRI 62.2 (Aug. 15, 2012), pp. 1241–1248. ISSN: 1053-8119 (page 26).
- [El +19] Loubna El Gueddari et al. "Calibrationless Oscar-Based Image Reconstruction in Compressed Sensing Parallel MRI". In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). IEEE. 2019, pp. 1532–1536 (page 44).
- [Ela+21] Jennifer Stine Elam et al. "The Human Connectome Project: A Retrospective". In: *NeuroImage* 244 (Dec. 1, 2021), p. 118543. ISSN: 1053-8119 (pages 2, 6).
- [Ell+20] Cameron T. Ellis et al. "Facilitating Open-Science with Realistic fMRI Simulation: Validation and Application". In: *PeerJ* 8 (Feb. 19, 2020), e8564. ISSN: 2167-8359 (pages 94, 95).
- [Eng+18] Maria Engel et al. "Single-Shot Spiral Imaging at 7 T". In: Magnetic Resonance in Medicine 80.5 (2018), pp. 1836–1846. ISSN: 1522-2594 (page 38).
- [Erh+12] Erik B. Erhardt et al. "SimTB, a Simulation Toolbox for fMRI Data under a Model of Spatiotemporal Separability". In: *NeuroImage* 59.4 (Feb. 15, 2012), pp. 4160–4167. ISSN: 1053-8119 (page 95).
- [Est+19] Oscar Esteban et al. "fMRIPrep: A Robust Preprocessing Pipeline for Functional MRI". In: Nature Methods 16.1 (1 Jan. 2019), pp. 111–116. ISSN: 1548-7105 (pages 29, 80).
- [Far+20] S Farrens et al. "PySAP: Python Sparse Data Analysis Package for Multidisciplinary Image Processing". In: Astronomy and Computing 32 (2020), p. 100402 (pages 41, 101, 102, 106, 136).
- [Fat+22] Ketan Fatania et al. A Plug-and-Play Approach to Multiparametric Quantitative MRI: Image Reconstruction Using Pre-Trained Deep Denoisers. Feb. 10, 2022.
 URL: http://arxiv.org/abs/2202.05269 (page 122).
- [Fei+23] David A. Feinberg et al. "Next-Generation MRI Scanner Designed for Ultra-High-Resolution Human Brain Imaging at 7 Tesla". In: *Nature Methods* 20.12 (Dec. 2023), pp. 2048–2057. ISSN: 1548-7105 (page 35).
- [Fes+05] J.A. Fessler et al. "Toeplitz-Based Iterative Image Reconstruction for MRI with Correction for Magnetic Field Inhomogeneity". In: *IEEE Transactions on Signal Processing* 53.9 (Sept. 2005), pp. 3393–3402. ISSN: 1941-0476 (page 72).
- [Fri07] K. J. Friston, ed. Statistical Parametric Mapping: The Analysis of Funtional Brain Images. 1st ed. Amsterdam ; Boston: Elsevier/Academic Press, 2007.
 647 pp. ISBN: 978-0-12-372560-8 (pages 29, 80).
- [Fri+00] K. J. Friston et al. "Nonlinear Responses in fMRI: The Balloon Model, Volterra Kernels, and Other Hemodynamics". In: *NeuroImage* 12.4 (Oct. 1, 2000), pp. 466–477. ISSN: 1053-8119 (page 16).
- [FSL17] FSL. "FSL Course Intro". 2017 (page 29).
- [Gao+25] Harry Gao et al. A Self-supervised Diffusion Bridge for MRI Reconstruction. Jan. 6, 2025. URL: http://arxiv.org/abs/2501.03430 (visited on 01/13/2025). Pre-published (page 64).

[GD14]	Matan Gavish and David L. Donoho. <i>The Optimal Hard Threshold for Singular Values Is</i> $4/Sqrt(3)$. June 4, 2014. URL: http://arxiv.org/abs/1305.5870. Pre-published (pages 3, 7, 80, 82).
[GD17]	Matan Gavish and David L. Donoho. "Optimal Shrinkage of Singular Values". In: <i>IEEE Transactions on Information Theory</i> 63.4 (Apr. 2017), pp. 2137–2152. ISSN: 1557-9654 (pages 3, 7, 80, 82, 83).
[Gil+23]	Chaithya Giliyar Radhakrishna et al. "Improving Spreading Projection Algorithm for Rapid K-Space Sampling Trajectories through Minimized off-Resonance Effects and Gridding of Low Frequencies". In: <i>Magnetic Resonance in Medicine</i> 90.3 (2023), pp. 1069–1085. ISSN: 1522-2594 (page 114).
[Glo12]	Gary H Glover. "Spiral Imaging in fMRI". In: (2012), p. 7 (pages 38, 39).
[Glo99]	Gary H. Glover. "Deconvolution of Impulse Response in Event-Related BOLD fMRI1". In: <i>NeuroImage</i> 9.4 (Apr. 1, 1999), pp. 416–429. ISSN: 1053-8119 (pages 17, 98).
[Gol+14]	Tom Goldstein et al. "Fast Alternating Direction Optimization Methods". In: <i>SIAM Journal on Imaging Sciences</i> 7.3 (Jan. 2014), pp. 1588–1623. ISSN: 1936-4954 (page 155).
[Gop+24]	Karthik Gopinath et al. "Synthetic Data in Generalizable, Learning-Based Neuroimaging". In: <i>Imaging Neuroscience</i> (Oct. 2024). ISSN: 2837-6056 (pages 93, 118).
[Gor+11]	Krzysztof Gorgolewski et al. "Nipype: A Flexible, Lightweight and Extensi- ble Neuroimaging Data Processing Framework in Python". In: <i>Frontiers in</i> <i>Neuroinformatics</i> 5 (2011). ISSN: 1662-5196 (page 80).
[Gor+16]	Krzysztof J. Gorgolewski et al. "The Brain Imaging Data Structure, a Format for Organizing and Describing Outputs of Neuroimaging Experiments". In: <i>Scientific Data</i> 3.1 (1 June 21, 2016), p. 160044. ISSN: 2052-4463 (page 29).
[GMC22]	Nadine N. Graedel, Karla L. Miller, and Mark Chiew. "Ultrahigh Resolution fMRI at 7T Using Radial-Cartesian TURBINE Sampling". In: <i>Magnetic Resonance in Medicine</i> 88.5 (2022), pp. 2058–2073. ISSN: 1522-2594 (pages 38, 40, 45, 96, 117).
[Gra+17a]	Nadine N. Graedel et al. "Motion Correction for Functional MRI with Three- Dimensional Hybrid Radial-Cartesian EPI". In: <i>Magnetic Resonance in Medicine</i> 78.2 (2017), pp. 527–540. ISSN: 1522-2594 (page 40).
[Gra+17b]	Vincent Gras et al. "Universal Pulses: A New Concept for Calibration-Free Parallel Transmission". In: <i>Magnetic Resonance in Medicine</i> 77.2 (Feb. 2017), pp. 635–643. ISSN: 1522-2594 (pages 25, 31).
[Gri+16]	Ludovica Griffanti et al. "Challenges in the reproducibility of clinical studies with resting state fMRI: An example in early Parkinson's disease". In: <i>Neuroimage</i> 124 (2016), pp. 704–713 (page 93).
[Gri+02]	Mark A. Griswold et al. "Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA)". In: <i>Magnetic Resonance in Medicine</i> 47.6 (2002), pp. 1202–1210. ISSN: 1522-2594 (pages 36, 115).
[GP95]	HáKon Gudbjartsson and Samuel Patz. "The Rician Distribution of Noisy Mri Data". In: <i>Magnetic Resonance in Medicine</i> 34.6 (1995), pp. 910–914. ISSN: 1522-2594 (page 80).

[Gül+23]	Merve Gülle et al. "Highly-Accelerated High-Resolution Multi-Echo fMRI Us- ing Self-Supervised Physics-Driven Deep Learning Reconstruction". In: 2023 IEEE 9th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP). 2023 IEEE 9th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP). Dec. 2023, pp. 196–200 (page 27).
[Ham+18]	Kerstin Hammernik et al. "Learning a variational network for reconstruction of accelerated MRI data". In: <i>Magnetic resonance in medicine</i> 79.6 (2018), pp. 3055–3071 (pages 60, 62, 122).
[HU20]	Martin Havlicek and Kâmil Uludağ. "A Dynamical Model of the Laminar BOLD Response". In: <i>NeuroImage</i> 204 (Jan. 1, 2020), p. 116209. ISSN: 1053-8119 (page 21).
[Hec+24]	Reinhard Heckel et al. "Deep Learning for Accelerated and Robust MRI Reconstruction". In: <i>Magnetic Resonance Materials in Physics, Biology and Medicine</i> 37.3 (July 1, 2024), pp. 335–368. ISSN: 1352-8661 (page 62).
[Hon+24a]	Tao Hong et al. Provable Preconditioned Plug-and-Play Approach for Compressed Sensing MRI Reconstruction. May 6, 2024. URL: http://arxiv.org/abs/2405. 03854. Pre-published (pages 122–126, 128).
[Hon+24b]	Tao Hong et al. "Provable preconditioned plug-and-play approach for compressed sensing MRI reconstruction". In: <i>IEEE Transactions on Computational Imaging</i> (2024) (page 135).
[Hua+21]	Tao Huang et al. "Neighbor2neighbor: Self-supervised denoising from single noisy images". In: <i>Proceedings of the IEEE/CVF conference on computer vision and pattern recognition</i> . 2021, pp. 14781–14790 (pages 59, 124).
[Hua+14]	Yue Huang et al. "Bayesian Nonparametric Dictionary Learning for Compressed Sensing MRI". In: <i>IEEE Transactions on Image Processing</i> 23.12 (Dec. 2014), pp. 5007–5019. ISSN: 1941-0042 (page 44).
[HLP22]	Samuel Hurault, Arthur Leclaire, and Nicolas Papadakis. "Gradient Step De- noiser for convergent Plug-and-Play". In: <i>International Conference on Learning</i> <i>Representations.</i> 2022 (pages 122, 123, 126).
[Iqb+20]	Asif Iqbal et al. "Adaptive Complex-Valued Dictionary Learning: Application to fMRI Data Analysis". In: <i>Signal Processing</i> 166 (Jan. 2020), p. 107263. ISSN: 01651684 (page 44).
[Jam+21]	Redouane Jamil et al. "Temporal SNR Optimization through RF Coil Combination in fMRI: The More, the Better?" In: <i>PLoS ONE</i> 16.11 (Nov. 8, 2021), e_592_0259. ISSN: 1932-6203 (page 114).
[Jin+06]	Tao Jin et al. "Source of Nonlinearity in Echo-Time-Dependent BOLD fMRI". In: <i>Magnetic Resonance in Medicine</i> 55.6 (2006), pp. 1281–1290. ISSN: 1522-2594 (pages 20, 98).
[Joc+06]	Thies H. Jochimsen et al. "Efficient Simulation of Magnetic Resonance Imaging with Bloch–Torrey Equations Using Intra-Voxel Magnetization Gradients". In: <i>Journal of Magnetic Resonance</i> 180.1 (May 2006), pp. 29–38. ISSN: 10907807 (page 95).
[JS14]	Julie Josse and Sylvain Sardy. Adaptive Shrinkage of Singular Values. Nov. 22, 2014. URL: http://arxiv.org/abs/1310.6602. preprint (pages 80, 82, 83).

[Jun+09]	Hong Jung et al. "K-t FOCUSS: A General Compressed Sensing Framework for High Resolution Dynamic MRI". In: <i>Magnetic Resonance in Medicine</i> 61.1 (Jan. 2009), pp. 103–116. ISSN: 1522-2594 (page 51).
[Kas+15]	Lars Kasper et al. "Monitoring, Analysis, and Correction of Magnetic Field Fluctuations in Echo Planar Imaging Time Series". In: <i>Magnetic Resonance in</i> <i>Medicine</i> 74.2 (2015), pp. 396–409. ISSN: 1522-2594 (page 72).
[Kas+22]	Lars Kasper et al. "Advances in Spiral fMRI: A High-Resolution Study with Single-Shot Acquisition". In: <i>NeuroImage</i> 246 (Feb. 2022), p. 118738. ISSN: 1053-8119 (page 38).
[KF18]	Donghwan Kim and Jeffrey A. Fessler. "Adaptive Restart of the Optimized Gradient Method for Convex Optimization". In: <i>Journal of Optimization Theory</i> and Applications 178.1 (July 2018), pp. 240–263. ISSN: 0022-3239, 1573-2878 (pages 106, 153).
[KO12]	Seong-Gi Kim and Seiji Ogawa. "Biophysical and Physiological Origins of Blood Oxygenation Level-Dependent fMRI Signals". In: Journal of Cerebral Blood Flow and Metabolism: Official Journal of the International Society of Cerebral Blood Flow and Metabolism 32.7 (July 2012), pp. 1188–1206. ISSN: 1559-7016 (pages 20, 117).
[Kno+11]	Florian Knoll et al. "Adapted Random Sampling Patterns for Accelerated MRI". In: <i>Magnetic Resonance Materials in Physics, Biology and Medicine</i> 24.1 (Feb. 1, 2011), pp. 43–50. ISSN: 1352-8661 (page 43).
[Kno+14]	Florian Knoll et al. "gpuNUFFT - An Open Source GPU Library for 3D Regrid- ding with Direct Matlab Interface". In: 2014 (page 70).
[Kon+16]	Jakub Konečný et al. "Mini-Batch Semi-Stochastic Gradient Descent in the Proximal Setting". In: <i>IEEE Journal of Selected Topics in Signal Processing</i> 10.2 (Mar. 2016), pp. 242–255. ISSN: 1941-0484 (page 51).
[KK17]	Ryoichi Kose and Katsumi Kose. "BlochSolver: A GPU-optimized Fast 3D MRI Simulator for Experimentally Compatible Pulse Sequences". In: <i>Journal of Magnetic Resonance</i> 281 (Aug. 1, 2017), pp. 51–65. ISSN: 1090-7807 (page 95).
[KBJ19]	Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2Void - Learning Denoising from Single Noisy Images. Apr. 5, 2019. URL: http://arxiv.org/abs/1811.10980 (visited on 04/25/2023). Pre-published (page 59).
[Kum+21]	Rajat Kumar et al. "Ground-Truth "Resting-State" Signal Provides Data-Driven Estimation and Correction for Scanner Distortion of fMRI Time-Series Dynamics". In: <i>NeuroImage</i> 227 (Feb. 15, 2021), p. 117584. ISSN: 1053-8119 (page 93).
[Kun+17]	Prantik Kundu et al. "Multi-Echo fMRI: A Review of Applications in fMRI Denoising and Analysis of BOLD Signals". In: <i>NeuroImage</i> . Cleaning up the fMRI Time Series: Mitigating Noise with Advanced Acquisition and Correction Strategies 154 (July 1, 2017), pp. 59–80. ISSN: 1053-8119 (page 27).
[Kwo+92]	K. K. Kwong et al. "Dynamic Magnetic Resonance Imaging of Human Brain Activity during Primary Sensory Stimulation". In: <i>Proceedings of the National</i> <i>Academy of Sciences of the United States of America</i> 89.12 (June 15, 1992), pp. 5675–5679. ISSN: 0027-8424 (pages 1, 5).
[Laj+22]	Hélène Lajous et al. "A Fetal Brain Magnetic Resonance Acquisition Numerical Phantom (FaBiAN)". In: <i>Scientific Reports</i> 12.1 (May 23, 2022), p. 8682. ISSN: 2045-2322 (page 95).

- [Lam+13] Fan Lam et al. "Accelerated fMRI Using Low-Rank Model and Sparsity Constraints". In: *Proc. Intl. Soc. Mag. Reson. Med.* 21 (May 2013) (page 54).
- [Lau+22] Rémi Laumont et al. "Bayesian Imaging Using Plug & Play Priors: When Langevin Meets Tweedie". In: SIAM Journal on Imaging Sciences 15.2 (June 2022), pp. 701–737. ISSN: 1936-4954 (page 62).
- [Law+19] Samuel JD Lawrence et al. "Laminar fMRI: Applications for cognitive neuroscience". In: *Neuroimage* 197 (2019), pp. 785–791 (pages 2, 6).
- [Laz+19] Carole Lazarus et al. "SPARKLING: Variable-Density k-Space Filling Curves for Accelerated T2*-Weighted MRI". In: Magnetic resonance in medicine 81.6 (2019), pp. 3643–3661 (pages 43, 122).
- [LSte+19] Caroline Le Ster et al. "Comparison of SMS-EPI and 3D-EPI at 7T in an fMRI Localizer Study with Matched Spatiotemporal Resolution and Homogenized Excitation Profiles". In: *PloS One* 14.11 (2019), e_286_0225. ISSN: 1932-6203 (pages 26, 37).
- [Leh+18] Jaakko Lehtinen et al. Noise2Noise: Learning Image Restoration without Clean Data. Oct. 29, 2018. URL: http://arxiv.org/abs/1803.04189 (visited on 12/17/2024). Pre-published (page 59).
- [LF19] Claire Yilin Lin and Jeffrey A. Fessler. "Efficient Dynamic Parallel MRI Reconstruction for the Low-Rank Plus Sparse Model". In: *IEEE Transactions on Computational Imaging* 5.1 (Mar. 2019), pp. 17–26. ISSN: 2333-9403 (pages 51, 105).
- [LH22] Kang Lin and Reinhard Heckel. "Vision Transformers Enable Fast and Robust Accelerated MRI". In: Proceedings of The 5th International Conference on Medical Imaging with Deep Learning. International Conference on Medical Imaging with Deep Learning. PMLR, Dec. 4, 2022, pp. 774–795 (page 59).
- [Lio+06] Michelle Liou et al. "A method for generating reproducible evidence in fMRI studies". In: *NeuroImage* 29.2 (2006), pp. 383–395 (page 93).
- [Liu+17] Fang Liu et al. "Fast Realistic MRI Simulations Based on Generalized Multi-Pool Exchange Tissue Model". In: *IEEE Transactions on Medical Imaging* 36.2 (Feb. 2017), pp. 527–537. ISSN: 1558-254X (page 95).
- [Liu16] Thomas T. Liu. "Noise Contributions to the fMRI Signal: An Overview". In: *NeuroImage* 143 (Dec. 1, 2016), pp. 141–151. ISSN: 1053-8119 (page 80).
- [Lus+08] M. Lustig et al. "Compressed Sensing MRI". In: *IEEE Signal Processing Magazine* 25.2 (Mar. 2008), pp. 72–82. ISSN: 1053-5888 (pages 40, 43).
- [LDP07] Michael Lustig, David Donoho, and John M Pauly. "Sparse MRI: The application of compressed sensing for rapid MR imaging". In: Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine 58.6 (2007), pp. 1182–1195 (page 122).
- [Mal09] S. G. Mallat. A Wavelet Tour of Signal Processing: The Sparse Way. 3rd ed. Amsterdam; Boston: Elsevier/Academic Press, 2009. 805 pp. ISBN: 978-0-12-374370-1 (page 41).
- [Man+13] José V. Manjón et al. "Diffusion Weighted Image Denoising Using Overcomplete Local PCA". In: *PLOS ONE* 8.9 (Sept. 3, 2013), e_1_7302. ISSN: 1932-6203 (pages 80, 82).

[MP67]	V. A. Marčenko and L. A. Pastur. "DISTRIBUTION OF EIGENVALUES FOR SOME SETS OF RANDOM MATRICES". In: <i>Mathematics of the USSR-Sbornik</i> 1.4 (Apr. 30, 1967), p. 457. ISSN: 0025-5734 (page 83).
[Mar+22]	Scott Marek et al. "Reproducible Brain-Wide Association Studies Require Thousands of Individuals". In: <i>Nature</i> 603.7902 (7902 Mar. 2022), pp. 654–660. ISSN: 1476-4687 (page 93).
[Mas+21]	Harry T. Mason et al. "Subspace-Constrained Approaches to Low-Rank fMRI Acceleration". In: <i>NeuroImage</i> 238 (Sept. 1, 2021), p. 118235. ISSN: 1053-8119 (page 53).
[Maz+23]	Audrey Mazancieux et al. "Brainstem fMRI Signaling of Surprise across Different Types of Deviant Stimuli". In: <i>Cell Reports</i> 42.11 (Nov. 28, 2023). ISSN: 2211-1247 (pages 2, 6).
[Moe+21]	Steen Moeller et al. "NOise Reduction with DIstribution Corrected (NORDIC) PCA in dMRI with Complex-Valued Parameter-Free Locally Low-Rank Processing". In: <i>NeuroImage</i> 226 (Feb. 1, 2021), p. 117539. ISSN: 1053-8119 (pages 3, 6, 80, 82, 87).
[Mor+22]	Thomas Moreau et al. "Benchopt: Reproducible, efficient and collaborative optimization benchmarks". In: Advances in Neural Information Process. Systems 35 (2022), pp. 25404–25421 (page 129).
[Muc+20]	M. J. Muckley et al. "TorchKbNufft: A High-Level, Hardware-Agnostic Non- Uniform Fast Fourier Transform". In: <i>ISMRM Workshop on Data Sampling & Im-</i> <i>age Reconstruction</i> . Source code available at https://github.com/mmuckley/torchkbnufft. 2020 (page 70).
[Muc+21]	Matthew J Muckley et al. "Results of the 2020 fastMRI challenge for machine learning MR image reconstruction". In: <i>IEEE transactions on medical imaging</i> 40.9 (2021), pp. 2306–2317 (page 122).
[Nak+23]	Johan Nakuci et al. "Within-Subject Reproducibility Varies in Multi-Modal, Longitudinal Brain Networks". In: <i>Scientific Reports</i> 13.1 (1 Apr. 24, 2023), p. 6699. ISSN: 2045-2322 (page 93).
[Nee19]	Derek Evan Nee. "fMRI replicability depends upon sufficient individual-level data". In: <i>Communications biology</i> 2.1 (2019), p. 130 (page 93).
[NNS22]	Rafael Neto Henriques, Sune Nørhøj Jespersen, and Noam Shemesh. "Hybrid PCA Denoising - Improving PCA Denoising in the Presence of Spatial Correlations". In: <i>Proceedings of the 30'th Scientific Meeting of the International Society for Magnetic Resonance in Medicine</i> (2022) (pages 80, 82, 83).
[Nil]	Nilearn. Nilearn (page 31).
[Nis96]	Dwight George Nishimura. <i>Principles of Magnetic Resonance Imaging.</i> Stanford University, 1996. 232 pp. (page 16).
[Nit14]	Atsushi Nitanda. "Stochastic Proximal Gradient Descent with Acceleration Techniques". In: <i>Proceedings of the 28th International Conference on Neural</i> <i>Information Processing Systems - Volume 1</i> . Vol. 1. NIPS'14. MIT Press, Dec. 8, 2014, pp. 1574–1582 (page 51).
[Oga+90]	S Ogawa et al. "Brain Magnetic Resonance Imaging with Contrast Dependent on Blood Oxygenation." In: <i>Proceedings of the National Academy of Sciences</i> 87.24 (Dec. 1990), pp. 9868–9872 (pages 1, 5).

- [OUL20] Frank Ong, Martin Uecker, and Michael Lustig. "Accelerating Non-Cartesian MRI Reconstruction Convergence Using k-Space Preconditioning". In: *IEEE* transactions on medical imaging 39.5 (May 2020), pp. 1646–1654. ISSN: 0278-0062 (page 40).
- [OL19] Ong Frank and Lustig Michael. "SigPy: A Python Package for High Performance Iterative Reconstruction". In: *ISMRM 2019*. ISMRM. 2019 (pages 70, 103).
- [OCS15] Ricardo Otazo, Emmanuel Candès, and Daniel K. Sodickson. "Low-Rank plus Sparse Matrix Decomposition for Accelerated Dynamic MRI with Separation of Background and Dynamic Components". In: *Magnetic Resonance in Medicine* 73.3 (2015), pp. 1125–1136. ISSN: 1522-2594 (pages 51, 105, 107).
- [Par+14] Dennis L Parker et al. "Phase reconstruction from multiple coil data using a virtual reference coil". In: Magnetic resonance in medicine 72.2 (2014), pp. 563– 569 (pages 124, 125, 136).
- [Ped21] Larson Peder. *MRI-education-resources and Principles of MRI E-book*. Version v1.0. Zenodo, Oct. 3, 2021 (page 16).
- [Pes+21] Jean-Christophe Pesquet et al. "Learning maximally monotone operators for image recovery". In: SIAM Journal on Imaging Sciences 14.3 (2021), pp. 1206– 1237 (pages 122, 123, 126).
- [PHA17] Andrii Y. Petrov, Michael Herbst, and V. Andrew Stenger. "Improving Temporal Resolution in fMRI Using a 3D Spiral Acquisition and Low Rank plus Sparse (L+S) Reconstruction". In: *NeuroImage* 157 (Aug. 2017), pp. 660–674. ISSN: 10538119 (pages 38, 39, 45, 51, 105, 107, 117).
- [Pfa+23] Laura Pfaff et al. "Self-supervised MRI denoising: leveraging Stein's unbiased risk estimator and spatially resolved noise maps". In: Scientific Reports 13.1 (2023), p. 22629 (page 124).
- [Pic+11] Davide Piccini et al. "Spiral Phyllotaxis: The Natural Way to Construct a 3D Radial Trajectory in MRI". In: *Magnetic Resonance in Medicine* 66.4 (2011), pp. 1049–1056. ISSN: 1522-2594 (page 38).
- [Pin+23] Walter H. L. Pinaya et al. Generative AI for Medical Imaging: Extending the MONAI Framework. July 27, 2023. URL: http://arxiv.org/abs/2307.15208. preprint (page 93).
- [Pin+20] Ana Luìsa Pinho et al. "Individual Brain Charting Dataset Extension, Second Release of High-Resolution fMRI Data for Cognitive Mapping". In: Scientific Data 7 (Oct. 16, 2020), p. 353 (pages 2, 6, 93).
- [PSS16] Rolf Pohmann, Oliver Speck, and Klaus Scheffler. "Signal-to-Noise Ratio and MR Tissue Parameters in Human Brain Imaging at 3, 7, and 9.4 Tesla Using Current Receive Coil Arrays". In: *Magnetic Resonance in Medicine* 75.2 (2016), pp. 801–809. ISSN: 1522-2594 (pages 25, 26).
- [Pol+18] Jonathan R Polimeni et al. "Analysis strategies for high-resolution UHF-fMRI data". In: *Neuroimage* 168 (2018), pp. 296–320 (page 132).
- [Pos+10] B. A. Poser et al. "Three Dimensional Echo-Planar Imaging at 7 Tesla". In: *NeuroImage* 51.1 (May 15, 2010), pp. 261–266. ISSN: 1053-8119 (pages 26, 96, 104).

[PMN10]	Benedikt A. Poser, Emily van Mierlo, and David G. Norris. "Exploring the Post-stimulus Undershoot with Spin-echo fMRI: Implications for Models of Neurovascular Response". In: <i>Human Brain Mapping</i> 32.1 (Dec. 14, 2010), p. 141 (page 24).
[PZJ23]	Aniket Pramanik, M Bridget Zimmerman, and Mathews Jacob. "Memory-efficient model-based deep learning with convergence and robustness guarantees". In: <i>IEEE transactions on computational imaging</i> 9 (2023), pp. 260–275 (page 122).
[Pru+99]	K. P. Pruessmann et al. "SENSE: Sensitivity Encoding for Fast MRI". In: Magnetic Resonance in Medicine 42.5 (Nov. 1999), pp. 952–962. ISSN: 0740-3194 (page 37).
[22]	Public Protocols. The Human Brain Project - brain charting task, Aug. 20, 2022 (page 83).
[PCP11]	N. Pustelnik, C. Chaux, and J. Pesquet. "Parallel Proximal Algorithm for Image Restoration Using Hybrid Regularization". In: <i>IEEE Transactions on</i> <i>Image Processing</i> 20.9 (Sept. 2011), pp. 2450–2462. ISSN: 1057-7149, 1941-0042 (page 156).
[QNJ18]	Tran Minh Quan, Thanh Nguyen-Duc, and Won-Ki Jeong. "Compressed Sensing MRI Reconstruction Using a Generative Adversarial Network With a Cyclic Loss". In: <i>IEEE Transactions on Medical Imaging</i> 37.6 (June 2018), pp. 1488–1497. ISSN: 1558-254X (page 59).
[Rad23]	Chaithya Giliyar Radhakrishna. "Designing and Learning Non-Cartesian k-Space Sampling Trajectories for Accelerated 3D MRI". PhD thesis. Université Paris- Saclay, Apr. 25, 2023 (page 16).
[RC23]	Chaithya Giliyar Radhakrishna and Philippe Ciuciu. "Jointly Learning Non-Cartesian k-Space Trajectories and Reconstruction Networks for 2D and 3D MR Imaging through Projection". In: <i>Bioengineering</i> 10.2 (2 Feb. 2023), p. 158. ISSN: 2306-5354 (pages 2, 6, 43, 144).
[Rad+22]	Chaithya Giliyar Radhakrishna et al. "MORE-SPARKLING: Non-Cartesian Trajectories with Minimized off-Resonance Effects". In: Joint Annual Meeting ISMRM-ESMRMB & ISMRT 31st Annual Meeting. May 7, 2022 (page 43).
[RCS21]	Zaccharie Ramzi, Philippe Ciuciu, and Jean-Luc Starck. XPDNet for MRI Reconstruction: An Application to the 2020 fastMRI Challenge. July 7, 2021. URL: http://arxiv.org/abs/2010.07290 (pages 60, 62).
[Ram+22]	Zaccharie Ramzi et al. "NC-PDNet: A density-compensated unrolled network for 2D and 3D non-Cartesian MRI reconstruction". In: <i>IEEE Transactions on</i> <i>Medical Imaging</i> 41.7 (2022), pp. 1625–1638 (pages 2, 6, 60, 122, 128).
[Rau10]	Holger Rauhut. "Compressive Sensing and Structured Random Matrices". In: <i>Theoretical Foundations and Numerical Methods for Sparse Recovery</i> . Ed. by Massimo Fornasier. DE GRUYTER, July 16, 2010, pp. 1–92. ISBN: 978-3-11-022614-0 (page 42).
[RCF17]	 Saiprasad Ravishankar, Il Yong Chun, and Jeffrey A. Fessler. "Physics-Driven Deep Training of Dictionary-Based Algorithms for MR Image Reconstruction". In: 2017 51st Asilomar Conference on Signals, Systems, and Computers. 2017 51st Asilomar Conference on Signals, Systems, and Computers. Oct. 2017, pp. 1859–1863 (page 44).

- [Rav+18] Saiprasad Ravishankar et al. "Deep Dictionary-Transform Learning for Image Reconstruction". In: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018). 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018). Apr. 2018, pp. 1208–1212 (page 44).
- [RMS21] Christoph A. Rettenmeier, Danilo Maziero, and V. Andrew Stenger. "Three Dimensional Radial Echo Planar Imaging for Functional MRI". In: *Magnetic* resonance in medicine 87.1 (Aug. 19, 2021), p. 193 (pages 38, 40).
- [REM17] Yaniv Romano, Michael Elad, and Peyman Milanfar. The Little Engine That Could: Regularization by Denoising (RED). Sept. 3, 2017. URL: http://arxiv. org/abs/1611.02862 (visited on 12/17/2024). Pre-published (page 62).
- [RFB15] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. May 18, 2015. URL: http:// arxiv.org/abs/1505.04597 (visited on 12/12/2024) (page 59).
- [San+13] Paula Sanz Leon et al. "The Virtual Brain: A Simulator of Primate Brain Network Dynamics". In: Frontiers in Neuroinformatics 7 (2013). ISSN: 1662-5196 (pages 94, 95, 145).
- [Sca+22] Jonathan Scarlett et al. "Theoretical Perspectives on Deep Learning Methods in Inverse Problems". In: *IEEE Journal on Selected Areas in Information Theory* 3.3 (Sept. 2022), pp. 433–453. ISSN: 2641-8770 (page 58).
- [Sch+22] Michael Schirner et al. "Brain Simulation as a Cloud Service: The Virtual Brain on EBRAINS". In: *NeuroImage* 251 (May 1, 2022), p. 118973. ISSN: 1053-8119 (pages 94, 117).
- [Sch+18] Jo Schlemper et al. "A Deep Cascade of Convolutional Neural Networks for Dynamic MR Image Reconstruction". In: *IEEE Transactions on Medical Imaging* 37.2 (Feb. 2018), pp. 491–503. ISSN: 1558-254X (page 60).
- [SN10] Andrey Shabalin and Andrew Nobel. Reconstruction of a Low-rank Matrix in the Presence of Gaussian Noise. July 23, 2010. URL: http://arxiv.org/abs/ 1007.4148. preprint (pages 80, 82).
- [Shi+21] Yu-hsuan Shih et al. "cuFINUFFT: A Load-Balanced GPU Library for General-Purpose Nonuniform FFTs". Mar. 25, 2021 (page 70).
- [Shi+22] Efrat Shimron et al. "Implicit Data Crimes: Machine Learning Bias Arising from Misuse of Public Data". In: *Proceedings of the National Academy of Sciences* 119.13 (Mar. 29, 2022), e2117203119 (page 63).
- [Smi+04] Stephen M. Smith et al. "Advances in Functional and Structural MR Image Analysis and Implementation as FSL". In: *NeuroImage* 23 Suppl 1 (2004), S208– 219. ISSN: 1053-8119 (page 80).
- [Sou+18] Roberto Souza et al. "An Open, Multi-Vendor, Multi-Field-Strength Brain MR Dataset and Analysis of Publicly Available Skull Stripping Methods Agreement". In: *NeuroImage*. Segmenting the Brain 170 (Apr. 15, 2018), pp. 482–494. ISSN: 1053-8119 (pages 62, 136).
- [SMF15] J.-L. Starck, Fionn Murtagh, and Jalal M. Fadili. Sparse Image and Signal Processing: Wavelets and Related Geometric Multiscale Analysis. Second edition. New York, NY: Cambridge University Press, 2015. 428 pp. ISBN: 978-1-107-08806-1 (page 149).

[SB21]	Robert W. Stobbe and Christian Beaulieu. "Three-Dimensional Yarnball k-Space Acquisition for Accelerated MRI". In: <i>Magnetic Resonance in Medicine</i> 85.4 (2021), pp. 1840–1854. ISSN: 1522-2594 (page 38).
[Sto+10]	Tony Stocker et al. "High-performance computing MRI simulations". In: <i>Magnetic Resonance in Medicine</i> 64.1 (2010), pp. 186–193. ISSN: 1522-2594 (page 95).
[Sun+24]	Yu Sun et al. "Provable Probabilistic Imaging Using Score-Based Generative Priors". In: <i>IEEE Transactions on Computational Imaging</i> 10 (2024), pp. 1290–1305. ISSN: 2333-9403 (page 62).
[SNF03]	B.P. Sutton, D.C. Noll, and J.A. Fessler. "Fast, Iterative Image Reconstruction for MRI in the Presence of Field Inhomogeneities". In: <i>IEEE Transactions on Medical Imaging</i> 22.2 (Feb. 2003), pp. 178–188. ISSN: 1558-254X (pages 30, 97).
[Ter23]	Matthieu Terris. "Learning priors for scalable computational imaging algorithms, from theory to application in radio astronomy". 2023 (page 61).
[Ter+23]	Matthieu Terris et al. "Plug-and-play imaging with model uncertainty quantifica- tion in radio astronomy". In: <i>arXiv preprint arXiv:2312.07137</i> (2023) (page 124).
[Ter+24]	Matthieu Terris et al. "Equivariant plug-and-play image reconstruction". In: <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> . 2024, pp. 25255–25264 (page 122).
[Tib96]	Robert Tibshirani. "Regression Shrinkage and Selection via the Lasso". In: <i>Journal of the Royal Statistical Society. Series B (Methodological)</i> 58.1 (1996), pp. 267–288. ISSN: 0035-9246 (page 153).
[Tri+05]	C. Triantafyllou et al. "Comparison of Physiological Noise at 1.5 T, 3 T and 7 T and Optimization of fMRI Acquisition Parameters". In: <i>NeuroImage</i> 26.1 (May 15, 2005), pp. 243–250. ISSN: 1053-8119 (page 30).
[TPW11]	Christina Triantafyllou, Jonathan R. Polimeni, and Lawrence L. Wald. "Physiological Noise and Signal-to-Noise Ratio in fMRI with Multi-Channel Array Coils". In: <i>NeuroImage</i> 55.2 (Mar. 15, 2011), pp. 597–606. ISSN: 1053-8119 (pages 3, 6, 80).
[UOT15]	Martin Uecker, Frank Ong, and J Tamir. "Berkley advanced reconstruction toolbox". In: <i>Proc. Intl. Soc. Mag. Reson. Med. 23.</i> Toronto, 2015 (pages 70, 103).
[Uec+14]	Martin Uecker et al. "ESPIRiT–an Eigenvalue Approach to Autocalibrating Parallel MRI: Where SENSE Meets GRAPPA". In: <i>Magnetic Resonance in Medicine</i> 71.3 (2014), pp. 990–1001. ISSN: 1522-2594 (page 44).
[UUB15]	Kamil Uludag, Kamil Ugurbil, and Lawrence Berliner, eds. <i>fMRI: From Nuclear Spins to Brain Functions</i> . Vol. 30. Biological Magnetic Resonance. Boston, MA: Springer US, 2015. ISBN: 978-1-4899-7590-4 978-1-4899-7591-1 (page 16).
[UK23]	Kâmil Uludağ and Lars Kasper. "Chapter 22 - BOLD fMRI: Physiology and Acquisition Strategies". In: <i>Advances in Magnetic Resonance Technology and Applications</i> . Ed. by Karin Markenroth Bloch, Maxime Guye, and Benedikt A. Poser. Vol. 10. Ultra-High Field Neuro MRI. Academic Press, Jan. 1, 2023, pp. 351–369 (pages 16, 20, 24, 98).

- [UMU09] Kâmil Uludağ, Bernd Müller-Bierl, and Kâmil Uğurbil. "An Integrative Model for Neuronal Activity-Induced Signal Changes for Gradient and Spin Echo Functional Imaging". In: *NeuroImage* 48.1 (Oct. 15, 2009), pp. 150–165. ISSN: 1053-8119 (pages 16, 21).
- [UVL20] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. "Deep Image Prior".
 In: International Journal of Computer Vision 128.7 (July 2020), pp. 1867–1888.
 ISSN: 0920-5691, 1573-1405 (page 62).
- [vdZwa+09] Wietske van der Zwaag et al. "fMRI at 1.5, 3 and 7 T: Characterising BOLD Signal Changes". In: *NeuroImage* 47.4 (Oct. 1, 2009), pp. 1425–1434. ISSN: 1053-8119 (pages 21, 24, 25).
- [Van+13] Signe J. Vannesjo et al. "Gradient System Characterization by Impulse Response Measurements with a Dynamic Field Camera". In: *Magnetic Resonance in Medicine* 69.2 (2013), pp. 583–593. ISSN: 1522-2594 (page 72).
- [VBW13] Singanallur V. Venkatakrishnan, Charles A. Bouman, and Brendt Wohlberg.
 "Plug-and-Play Priors for Model Based Reconstruction". In: 2013 IEEE Global Conference on Signal and Information Processing. 2013 IEEE Global Conference on Signal and Information Processing (GlobalSIP). Austin, TX, USA: IEEE, Dec. 2013, pp. 945–948. ISBN: 978-1-4799-0248-4 (pages 62, 122, 123).
- [Ver+16] Jelle Veraart et al. "Denoising of Diffusion MRI Using Random Matrix Theory". In: *NeuroImage* 142 (Nov. 15, 2016), pp. 394–406. ISSN: 1095-9572 (pages 3, 6, 80, 82, 83).
- [VP21] Olivia Viessmann and Jonathan R. Polimeni. "High-Resolution fMRI at 7 Tesla: Challenges, Promises and Recent Developments for Individual-Focused fMRI Studies". In: *Current opinion in behavioral sciences* 40 (Aug. 2021), pp. 96–104. ISSN: 2352-1546 (pages 2, 6, 93, 132).
- [Vin+14] Thomas Vincent et al. "Flexible Multivariate Hemodynamics fMRI Data Analyses and Simulations with PyHRF". In: *Frontiers in Neuroscience* 8 (2014). ISSN: 1662-453X (pages 17, 117).
- [Viz+21] Luca Vizioli et al. "Lowering the Thermal Noise Barrier in Functional Brain Mapping with Magnetic Resonance Imaging". In: *Nature Communications* 12.1 (1 Aug. 30, 2021), p. 5181. ISSN: 2041-1723 (pages 3, 6, 80, 82, 83, 87, 141).
- [WF23] Guanhua Wang and Jeffrey A. Fessler. "Efficient Approximation of Jacobian Matrices Involving a Non-Uniform Fast Fourier Transform (NUFFT)". In: *IEEE Transactions on Computational Imaging* 9 (2023), pp. 43–54. ISSN: 2333-9403, 2334-0118, 2573-0436 (page 74).
- [Wan+22] Guanhua Wang et al. "B-Spline Parameterized Joint Optimization of Reconstruction and K-Space Trajectories (BJORK) for Accelerated 2D MRI". In: *IEEE Transactions on Medical Imaging* 41.9 (Sept. 2022), pp. 2318–2330. ISSN: 1558-254X (page 43).
- [Wan+23] Huan Wang et al. "White Matter BOLD Signals at 7 Tesla Reveal Visual Field Maps in Optic Radiation and Vertical Occipital Fasciculus". In: *NeuroImage* 269 (Apr. 1, 2023), p. 119916. ISSN: 1053-8119 (page 17).
- [Web16] Andrew Webb. "Radiofrequency Coils". In: (Apr. 27, 2016) (page 22).
- [Wei+21a] Kaixuan Wei et al. TFPnP: Tuning-free Plug-and-Play Proximal Algorithm with Applications to Inverse Imaging Problems. Version 3. Sept. 18, 2021. URL: http://arxiv.org/abs/2012.05703 (pages 122, 124).

[Wei+21b]	Tomer Weiss et al. <i>PILOT: Physics-Informed Learned Optimized Trajectories</i> for Accelerated MRI. Apr. 13, 2021. URL: http://arxiv.org/abs/1909.05773 (visited on 12/02/2024). Pre-published (page 43).
[WR14]	Marijke Welvaert and Yves Rosseel. "A Review of fMRI Simulation Studies". In: <i>PLOS ONE</i> 9.7 (July 21, 2014), e_53_1019. ISSN: 1932-6203 (pages 93, 94).
[Wel+11]	Marijke Welvaert et al. "neuRosim: An R Package for Generating fMRI Data". In: Journal of Statistical Software 44 (Oct. 31, 2011), pp. 1–18. ISSN: 1548-7660 (pages 94, 95).
[Wu+24]	Zihui Wu et al. Principled Probabilistic Imaging Using Diffusion Models as Plug-and-Play Priors. Nov. 6, 2024. URL: http://arxiv.org/abs/2405.18782 (visited on 12/18/2024). Pre-published (page 62).
[XWJ20]	Yaochen Xie, Zhengyang Wang, and Shuiwang Ji. Noise2Same: Optimizing A Self-Supervised Bound for Image Denoising. Oct. 22, 2020. URL: http://arxiv. org/abs/2010.11971 (visited on 12/17/2024). Pre-published (page 59).
[YHU08]	Essa Yacoub, Noam Harel, and Kâmil Uğurbil. "High-field fMRI unveils orienta- tion columns in humans". In: <i>Proceedings of the National Academy of Sciences</i> 105.30 (2008), pp. 10607–10612 (pages 2, 6).
[Yad19]	Omry Yadan. Hydra - A Framework for Elegantly Configuring Complex Applica- tions. Github. 2019 (page 103).
[Yam+20a]	Burhaneddin Yaman et al. "Self-Supervised Learning of Physics-Guided Recon- struction Neural Networks without Fully-Sampled Reference Data". In: <i>Magnetic</i> <i>Resonance in Medicine</i> 84.6 (Dec. 2020), pp. 3172–3191. ISSN: 0740-3194, 1522- 2594 (pages 59, 64).
[Yam+20b]	Burhaneddin Yaman et al. "Self-Supervised Physics-Based Deep Learning MRI Reconstruction Without Fully-Sampled Data". In: 2020 IEEE 17th Interna- tional Symposium on Biomedical Imaging (ISBI). 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI). Apr. 2020, pp. 921–925 (page 64).
[Yeo+11]	B. T. Thomas Yeo et al. "The Organization of the Human Cerebral Cortex Estimated by Intrinsic Functional Connectivity". In: <i>Journal of Neurophysiology</i> 106.3 (Sept. 2011), pp. 1125–1165. ISSN: 1522-1598 (page 117).
[Z A+24]	Z. Amor et al. "Achieving High Temporal Resolution Using a Sliding-Window Approach for SPARKLING fMRI Data: A Simulation Study". In: <i>ISMRM Annual Meeting</i> . Singapore, May 2024 (page 9).
[Zbo+18]	Jure Zbontar et al. "fastMRI: An open dataset and benchmarks for accelerated MRI". In: <i>arXiv preprint arXiv:1811.08839</i> (2018) (pages 62, 122, 124).
[Zha+23]	Minye Zhan et al. "Does the Visual Word Form Area Split in Bilingual Readers? A Millimeter-Scale 7-T fMRI Study". In: <i>Science Advances</i> 9.14 (Apr. 5, 2023), eadf6140 (pages 2, 6).
[Zha+21a]	Kai Zhang et al. "Plug-and-Play Image Restoration with Deep Denoiser Prior". In: <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> 44.10 (2021), pp. 6360–6376 (page 135).
[Zha+21b]	Kai Zhang et al. "Plug-and-play image restoration with deep denoiser prior". In: <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> 44.10 (2021), pp. 6360–6376 (pages 59, 60, 123, 125–127).

- [Zhu+18] Bo Zhu et al. "Image Reconstruction by Domain-Transform Manifold Learning". In: *Nature* 555.7697 (Mar. 2018), pp. 487–492. ISSN: 1476-4687 (page 59).
- [Zhu+22] Wei Zhu et al. "Denoise Functional Magnetic Resonance Imaging With Random Matrix Theory Based Principal Component Analysis". In: *IEEE Transactions* on Biomedical Engineering 69.11 (Nov. 2022), pp. 3377–3388. ISSN: 1558-2531 (page 83).
- [Zu+24] Zhongliang Zu et al. "The Missing Third Dimension–Functional Correlations of BOLD Signals Incorporating White Matter". In: *Science Advances* 10.4 (Jan. 26, 2024), eadi0616 (page 17).

UNIVERSITE PARIS-SACLAY

ÉCOLE DOCTORALE Physique et ingénierie: Electrons, Photons, Sciences du vivant (EOBE)

Titre: De l'échantillonnage compressé à l'apprentissage profond pour la reconstruction d'Imagerie par Résonance Magnétique fonctionnelle à très haut champ

Mots clés: neuroimagerie fonctionnelle; apprentissage profond; échantillonnage compressé. **Résumé:** L'imagerie par résonance magnétique fonctionnelle (IRMf) est une technique d'imagerie cérébrale largement utilisée, basée sur l'effet BOLD (Blood Oxygenation Level Dependent), un indicateur indirect de l'activité neuronale. L'IRMf est confrontée à un compromis entre la résolution spatiale, la résolution temporelle, le rapport signal sur bruit et la couverture cérébrale. Atteindre une couverture complète du cerveau avec une haute résolution spatiale et temporelle reste un défi.

Cette thèse de doctorat vise à améliorer les étapes d'acquisition et de reconstruction du pipeline IRMf afin d'atteindre une résolution plus élevée, en optimisant à terme les deux processus conjointement. À cette fin, nous avons développé « SNAKE », un simulateur open source d'IRMf, depuis l'activation neuronale et les réponses BOLD jusqu'à l'acquisition des données brutes. Ce simulateur permet la génération de données à haute résolution spatio-temporelle, facilitant le développement de méthodes d'acquisition et de reconstruction optimisées. En utilisant SNAKE, nous avons démontré l'efficacité d'une stratégie de démarrage à chaud pour la reconstruction IRMf avec des schémas d'acquisition dynamiques. Enfin, nous avons également exploré l'intégration de méthodes de débruitage dans la reconstruction en utilisant une méthode « Plug-and-Play ». Des résultats préliminaires ont été obtenus dans le contexte de l'IRM anatomique, ouvrant la voie à des travaux futurs pour l'étendre à l'IRMf.

Title: From Compressed Sensing to deep learning based methods for functional Magnetic Resonance Imaging reconstruction at ultra high field

Keywords: Functional MRI; Deep Learning; Compressed Sensing.

Abstract: Functional Magnetic Resonance Imaging (fMRI) is a widely used brain imaging technique based on the Blood Oxygenation Level Dependent (BOLD) effect, a proxy for neural activity. fMRI faces a trade-off between spatial resolution, temporal resolution, signal-to-noise ratio, and brain coverage. Achieving whole-brain coverage at high spatial and temporal resolution remains a challenge.

This PhD thesis aims to improve the acquisition and reconstruction stages of the fMRI pipeline to reach higher resolution, ultimately optimizing both processes jointly. To this end, we developed "SNAKE" an open-source fMRI simulator from neuronal activation and BOLD responses to raw data acquisition. This simulator enables the generation of high spatio-temporal resolution data, facilitating the development of optimized acquisition and reconstruction methods. Using SNAKE, we demonstrated the effectiveness of a warm-start strategy for fMRI reconstruction with dynamic acquisition patterns. Finally, we also explored integrating denoising methods into the reconstruction using a "Plug-and-Play" method. Preliminary results are obtained the context of anatomical MRI, setting the stage for future work extending it to fMRI.

Maison du doctorat de Université Paris-Saclay

^{2&}lt;sup>e</sup> étage, aile ouest, École normale supérieure Paris-Saclay

⁴ avenue des Sciencs

⁹¹¹⁹⁰ Gif-sur-Yvette, France