



HAL
open science

Off-the-grid data-driven optimization of sampling schemes in MRI

Alban Gossard, Frédéric de Gournay, Pierre Weiss

► **To cite this version:**

Alban Gossard, Frédéric de Gournay, Pierre Weiss. Off-the-grid data-driven optimization of sampling schemes in MRI. ITWIST 2020, Dec 2020, Nantes, France. hal-02945903

HAL Id: hal-02945903

<https://hal.science/hal-02945903v1>

Submitted on 2 Oct 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Off-the-grid data-driven optimization of sampling schemes in MRI

Alban Gossard^{1,2}, Frédéric de Gournay^{1,2} and Pierre Weiss^{1,2}.

¹Institut de Mathématiques de Toulouse. ²Institut des Technologies Avancées du Vivant. Université de Toulouse & INSA & CNRS.

Abstract— We propose a novel learning based algorithm to generate efficient and physically plausible sampling patterns in MRI. This method has a few advantages compared to recent learning based approaches: i) it works off-the-grid and ii) allows to handle arbitrary physical constraints. These two features allow for much more versatility in the sampling patterns that can take advantage of all the degrees of freedom offered by an MRI scanner. The method consists in a high dimensional optimization of a cost function defined implicitly by an algorithm. We propose various numerical tools to address this numerical challenge.

1 Introduction

The design of efficient sampling patterns in MRI is a critical issue with a long history [3] and a renewed interest in recent years with the advent of compressed sensing and deep learning.

State-of-the-art The most recent trends can be separated in two families

Compressed sensing theory: In a recent set of works, the theory of compressed sensing was improved to more closely fit the practical issues of MRI [6, 1, 5, 4]. In a nutshell, these works suggest that good sampling schemes should have a variable density: the low frequencies should be sampled more densely than the high frequencies (though this has little to do with the quantity of energy present in the signal) and the samples should cover the space locally uniformly. This led a few authors to generate sampling schemes by minimizing the distance between a measure μ belonging to a set of admissible measures and a continuous target probability density function π :

$$\inf_{\mu \in \mathcal{A}} \text{dist}(\mu, \pi), \quad (1)$$

where dist is a distance that metrizes the weak convergence, such as a discrepancy [10, 7] or the Wasserstein distance [16] and \mathcal{A} is a set of admissible probability measures such as the set of discrete measures supported on M points $\mathcal{M}_M = \{\mu = \frac{1}{M} \sum_{m=1}^M \delta_{x_m}, x_m \in \mathbb{R}^d\}$ or more exotic sets of constraints that more closely describe the physical constraints of a scanner. This approach led to remarkable practical results [15] that are currently evaluated for clinical routine.

Learning: Motivated by the recent breakthroughs of learning and deep learning, many authors recently tried to *learn* either the reconstructor [12], the sampling pattern [9, 19], or both [13]. In [9], the authors propose a greedy algorithm that generates a sampling pattern by iteratively selecting a discrete horizontal line that maximizes the SNR of the reconstructed image. A similar principle is proposed in [13], but there, the reconstructor is learnt simultaneously. In [19], the authors adopt a similar approach, but replace the greedy algorithm by a bi-level programming approach that controls the number of sampling points using an ℓ^1 penalization. Overall, all those works suffer from the same limitations:

- The sampling points are required to live on a Cartesian grid, which is suboptimal.
- The methods cannot incorporate advanced constraints on the sampling trajectory and therefore focus on “rigid” constraints such as imposing to sample horizontal lines.
- The methods are computationally intensive, which may be not be so critical since the sampling schemes are generated offline.

To the best of our knowledge, the only paper that addresses the above criticisms is the just posted [20], where the authors simultaneously optimize a reconstructor and a sampling scheme by performing a local optimization of a well initialized trajectory.

Our contribution In this paper, we propose to blend both approaches by using a data-driven distance in (1) rather than more principled approaches. This allows us to avoid all the above-mentioned flaws. This also leads us to implement a set of advanced numerical routines to address the computational challenges raised by the proposed cost function.

2 The proposed approach

2.1 Preliminaries

We assume that both a set of training images $\mathbf{x} = (x_1, \dots, x_K) \in \mathbb{C}^{N \times K}$ and a differentiable image quality metric $\eta : \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}_+$ are available. In this work, we will simply consider the squared ℓ^2 distance $\eta(\hat{x}, x) = \frac{1}{2} \|\hat{x} - x\|_2^2$. In what follows, we let $\xi \in (\mathbb{R}^d)^M$ denote a set of locations in the k -space (or Fourier domain), $y \in \mathbb{C}^M$ denote a set of k -space measurements and $R : \mathbb{C}^M \times (\mathbb{R}^d)^M \rightarrow \mathbb{C}^N$ denote a fixed reconstructor, i.e. for a sampling scheme $\xi \in (\mathbb{R}^d)^M$ and a measurement vector $y \in \mathbb{C}^M$, we let $\hat{x} = R(\xi, y)$ denote the reconstructed image. We let $A(\xi) \in \mathbb{C}^{M \times N}$ denote the forward Fourier transform defined for all $1 \leq m \leq M$ and $x \in \mathbb{C}^N$

$$[A(\xi)x]_m = \sum_{n=1}^N x_n \exp(-i\langle p_n, \xi_m \rangle), \quad (2)$$

where $p_n \in \{-N/2, \dots, N/2 - 1\}^d$ are the positions of grid points in the image space. For a regularization parameter $\lambda > 0$, we consider a Tikhonov reconstructor:

$$R_1(\xi, y) = \underset{x \in \mathbb{C}^N}{\text{argmin}} \frac{1}{2} \|A(\xi)x - y\|_2^2 + \frac{\lambda}{2} \|x\|_2^2, \quad (3)$$

and a nonlinear compressed sensing type reconstructor

$$R_2(\xi, y) = \Psi z^*, \quad z^* = \underset{z \in \mathbb{C}^P}{\text{argmin}} \frac{1}{2} \|A(\xi)\Psi z - y\|_2^2 + \lambda \|z\|_1, \quad (4)$$

where $\Psi \in \mathbb{C}^{N \times P}$ is a redundant wavelet transform.

2.2 The principle

The goal here is to replace the distance in (1) by a data-driven cost function. A natural choice reads:

$$\min_{\xi \in \Xi} F(\xi) := \mathbb{E} \left(\sum_{x \in \mathbf{x}} \eta(R(\xi, A(\xi)x + b), x) \right), \quad (\mathcal{P})$$

where $b \sim \mathcal{N}(0, \sigma^2 I_N)$ and $\Xi \subseteq (\mathbb{R}^d)^M$ describes the physical constraints. In words, the term $A(\xi)x + b$ represents noisy data acquisition that we want to reconstruct as well as possible, in average, using the reconstructor R . The expectation is taken w.r.t. the noise realizations.

Differentiating the reconstructors Solving (\mathcal{P}) is a real computational challenge. It is high dimensional, the cost function does not have a simple analytic formula and its regularity properties are unclear. Since the cost function is defined through another minimization problem, (\mathcal{P}) can be interpreted as a bi-level optimization problem. Various approaches are available to solve it [2]. Here we will follow the approach suggested in [18]. Instead of solving the lower-level minimization problems (3) or (4) exactly, we assume that they are solved approximately using iterative algorithms such as a conjugate gradient method or a proximal gradient descent. The main idea is then to differentiate the algorithm using dedicated libraries such as PyTorch instead of the minimizer itself.

Implementing and differentiating the NUFT The fast implementation of the linear mapping $A(\xi)$ is the backbone of our approach. It corresponds to the non uniform Fourier transform (NUFT). Various efficient approximate implementations have been devised over the past [8, 11, 14] and Python toolboxes begin to emerge [17]. Our experience using them however led to unstable results due to significant numerical errors. In this work, we therefore opted for a direct (naive) implementation of the NUFT on massively parallel architectures, following the numerical experiments conducted in KeOps. The main observation is that for a GPU with 1TFlop, applying the NUFT to small 128×128 images (which is typical in this field) just requires a fraction of second, which is compatible with large scale computations. We therefore implemented a homemade NUFT within PyTorch, allowing for automatic differentiation.

Optimizing the cost function The previous details allow to automatically compute the derivative of F w.r.t. ξ when replacing the expectation by an empirical average. This in turn allows to use any off-the-shelf optimization solver. In this preliminary work, we simply set $\Xi = (\mathbb{R}^d)^M$ (i.e. no constraints between samples), and $b = 0$ (no noise in the measurements) and used a limited memory BFGS algorithm. More advanced stochastic gradient approaches are expected to be used later.

3 Results

Here we report preliminary results with this approach. Two 64×64 images are studied and we compare 3 patterns subsampled at a factor 3.3. Our approach is abbreviated OSP (optimal sampling pattern), we also use a low-frequency pattern (LF) and a variable density sampler with a uniform density (VDS). The images are reconstructed both with the linear (3) and nonlinear reconstructors (4). Without surprise, LF is good at reconstructing global shape of images and removing noise but the VDS performs better to reconstruct details (with the disadvantage of




Figure 1: Original images (1a) square and (1b) phantom. Different k -space sampling pattern: (1d) VDS and (1c) LF random sampling.




Figure 2: Reconstructed images with Tikhonov reconstructor (3) for LF (2a), VDS (2b), OSP (2c) given in (2d).




Figure 3: Reconstructed images with nonlinear reconstructor (4) for LF (3a), VDS (3b), OSP (3c) given in (3d).




Figure 4: Reconstructed images with Tikhonov reconstructor (3) for LF (4a), VDS (4b), OSP (4c) given in (4d).




Figure 5: Reconstructed images with nonlinear reconstructor (4) for LF (5a), VDS (5b), OSP (5c) given in (5d).

generating noise). Our OSP combines both advantages. Our approach shows that choosing an optimized k -space improves the peak noise-to-signal ratio (PSNR) between 2dB and 10dB for these test-cases in comparison to variable density sampling or standard low-frequency sampling. In the future, we plan to focus on learning a reconstructor, adding physical constraints to the set Ξ and introducing the noise b .

Acknowledgments

This work was supported by the ANR Optimization on Measures Spaces and of ANR-3IA Artificial and Natural Intelligence Toulouse Institute.

References

- [1] Ben Adcock, Anders C Hansen, Clarice Poon, and Bogdan Roman. Breaking the coherence barrier: A new theory for compressed sensing. In *Forum of Mathematics, Sigma*, volume 5. Cambridge University Press, 2017.
- [2] Jonathan F Bard. *Practical bilevel optimization: algorithms and applications*, volume 30. Springer Science & Business Media, 2013.
- [3] Matt A Bernstein, Kevin F King, and Xiaohong Joe Zhou. *Handbook of MRI pulse sequences*. Elsevier, 2004.
- [4] Claire Boyer, Jérémie Bigot, and Pierre Weiss. Compressed sensing with structured sparsity and structured acquisition. *Applied and Computational Harmonic Analysis*, 46(2):312–350, 2019.
- [5] Claire Boyer, Nicolas Chauffert, Philippe Ciuciu, Jonas Kahn, and Pierre Weiss. On the generation of sampling schemes for magnetic resonance imaging. *SIAM Journal on Imaging Sciences*, 9(4):2039–2072, 2016.
- [6] Nicolas Chauffert, Philippe Ciuciu, Jonas Kahn, and Pierre Weiss. Variable density sampling with continuous trajectories. *SIAM Journal on Imaging Sciences*, 7(4):1962–1992, 2014.
- [7] Nicolas Chauffert, Philippe Ciuciu, Jonas Kahn, and Pierre Weiss. A projection method on measures sets. *Constructive Approximation*, 45(1):83–111, 2017.
- [8] Jeffrey A Fessler and Bradley P Sutton. Nonuniform fast fourier transforms using min-max interpolation. *IEEE transactions on signal processing*, 51(2):560–574, 2003.
- [9] Baran Gözcü, Rabeeh Karimi Mahabadi, Yen-Huan Li, Efe Ilıcak, Tolga Cukur, Jonathan Scarlett, and Volkan Cevher. Learning-based compressive mri. *IEEE transactions on medical imaging*, 37(6):1394–1406, 2018.
- [10] Gräf, Manuel and Potts, Daniel and Steidl, Gabriele. Quadrature errors, discrepancies, and their relations to halftoning on the torus and the sphere. *SIAM Journal on Scientific Computing*, 34(5):A2760–A2791, 2012.
- [11] Leslie Greengard and June-Yub Lee. Accelerating the nonuniform fast fourier transform. *SIAM review*, 46(3):443–454, 2004.
- [12] Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P Recht, Daniel K Sodickson, Thomas Pock, and Florian Knoll. Learning a variational network for reconstruction of accelerated mri data. *Magnetic resonance in medicine*, 79(6):3055–3071, 2018.
- [13] Kyong Hwan Jin, Michael Unser, and Kwang Moo Yi. Self-supervised deep active accelerated mri. *arXiv preprint arXiv:1901.04547*, 2019.
- [14] Jens Keiner, Stefan Kunis, and Daniel Potts. Using nfft 3—a software library for various nonequispaced fast fourier transforms. *ACM Transactions on Mathematical Software (TOMS)*, 36(4):1–30, 2009.
- [15] Carole Lazarus, Pierre Weiss, Nicolas Chauffert, Franck Mauconduit, Loubna El Gueddari, Christophe Destrieux, Ilyess Zemmoura, Alexandre Vignaud, and Philippe Ciuciu. Sparkling: variable-density k-space filling curves for accelerated t_2^* -weighted mri. *Magnetic resonance in medicine*, 81(6):3643–3661, 2019.
- [16] Lebrat, Léo and de Gournay, Frédéric and Kahn, Jonas and Weiss, Pierre. Optimal transport approximation of 2-dimensional measures. *SIAM Journal on Imaging Sciences*, 12(2):762–787, 2019.
- [17] Jyh-Miin Lin. Python non-uniform fast fourier transform (pynufft): An accelerated non-cartesian mri package on a heterogeneous platform (cpu/gpu). *Journal of Imaging*, 4(3):51, 2018.
- [18] Peter Ochs, René Ranftl, Thomas Brox, and Thomas Pock. Bilevel optimization with nonsmooth lower level problems. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 654–665. Springer, 2015.
- [19] Ferdia Sherry, Martin Benning, Juan Carlos De los Reyes, Martin J Graves, Georg Maierhofer, Guy Williams, Carola-Bibiane Schönlieb, and Matthias J Ehrhardt. Learning the sampling pattern for mri. *arXiv preprint arXiv:1906.08754*, 2019.
- [20] Tomer Weiss, Ortal Senouf, Sanketh Vedula, Oleg Michailovich, Michael Zibulevsky, and Alex Bronstein. Pilot: Physics-informed learned optimal trajectories for accelerated mri. *arXiv preprint arXiv:1909.05773*, 2020.