Compressed quantitative MRI: Bloch response recovery through iterated projection

COMPRESSED QUANTITATIVE MRI: BLOCH RESPONSE RECOVERY THROUGH ITERATED PROJECTION

Mike Davies[⋆⋆] *Gilles Puy*[†] *Pierre Vandergheynst*[†] *Yves Wiaux*[‡]

[⋆]Institute for Digital Communications (IDCom), The University of Edinburgh, EH9 3JL, UK.
[†]Institute of Electrical Engineering, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne, CH.
[‡]Institute of Sensors, Signals, and Systems, Heriot-Watt University, EH14 4AS, UK.

ABSTRACT

Inspired by the recently proposed Magnetic Resonance Fingerprinting technique, we develop a principled compressed sensing framework for quantitative MRI. The three key components are: a random pulse excitation sequence following the MRF technique; a random EPI subsampling strategy and an iterative projection algorithm that imposes consistency with the Bloch equations. We show that, as long as the excitation sequence possesses an appropriate form of persistent excitation, we are able to achieve accurate recovery of the proton density, T_1, T_2 and off-resonance maps simultaneously from a limited number of samples.

Index Terms— compressed sensing, MRI, Bloch equations, manifolds, Johnston-Linderstrauss embeddings.

1. INTRODUCTION

In the recent paper [1], a new type of MRI acquisition scheme called Magnetic Resonance Fingerprinting (MRF) is presented for the quantification of multiple tissue properties simultaneously through a single acquisition process. The procedure is composed of 4 key ingredients: (1) the material magnetization is excited through a sequence of random RF pulses; (2) after each pulse the response is recorded through measurements taken from a small portion of k-space; (3) a sequence of highly aliased magnetization response images are formed using back projection; and (4) parameter maps (proton density, ρ, T_1, T_2 and off-resonance, δf) are formed using a bank of matched filters comparing the "noisy" magnetization responses for each voxel with the predicted magnetization response for given parameter sets.

Inspired by this technique, we investigate this idea from a compressed sensing (CS) perspective. In [1], it was mentioned that MRF was itself inspired by the recent growth of compressed sensing techniques in MRI, however, the exact link to CS was not made explicit and the paper does not consider a full CS formulation. Indeed the role of sparsity, random excitation and sampling are not clarified.

[⋆]This work was initiated during a research visit by MD to EPFL funded by EPSRC grant EP/K032275/1.

Here we identify separate roles for the pulse excitation and the subsampling of k-space. We identify the Bloch response manifold as the appropriate low dimensional signal model on which the CS acquisition is performed and interpret the "model-based" dictionary of [1] as a natural discretization of this manifold. We then leverage recent results from [2] and develop a recovery algorithm with good theoretical guarantees. We conclude with some simulations to demonstrate the efficacy of our approach.

2. IR-SSFP EXCITATION

2.1. The Bloch response manifold

The MRF process is based upon an Inversion Recovery Steady State Free Precession (IR-SSFP) excitation sequence (see, *e.g.*, [3, 4]). Let $i = 1, \dots, N$, index the voxels of the imaged slice. We will assume that within each voxel a single isochromat dominates. The MRF excitation generates a magnetization field that can be observed at each excitation pulse. This field at voxel i at a given time t is a function of the excitation parameters at time t (namely, the flip angle α_t and the repetition time TR_t), the magnetization at time $t-1$, the overall magnetic field, the unknown parameters $\theta_i = \{T_1, T_2, \delta f\} \in \mathcal{M}$ associated with the local isochromat, and the voxel's proton density $\rho_i \geq 0$ [1]. This overall dynamics for an isochromat can be described by a parametrically excited linear system [3, 4].

Now and subsequently, we will denote the magnetization image sequence by a matrix $X \in \mathbb{C}^{N \times L}$, with $X_{i,t}$ denoting the magnetization for voxel i at the t^{th} read time. Given that the initial magnetization is known, the magnetization response at any voxel can be written as a parametric nonlinear mapping from $\{\rho_i, \theta_i\}$ to the sequence $X_{i,:}$ as follows

$$X_{i,:} = \rho_i B(\theta_i; \alpha, \mathrm{TR}) \in \mathbb{C}^{1 \times L}, \tag{1}$$

where L is the excitation sequence length, and $B : \mathcal{M} \to \mathbb{C}^{1 \times L}$ is a smooth mapping induced by the Bloch equation

[1]We treat the idealized case where ρ_i is assumed to be the actual proton density. It is common practice in MRI to use ρ to represent a combination of other factors such as coil sensitivity. In this case we simply let $\rho_i \in \mathbb{C}$.

dynamics. Note that we are representing the magnetization response sequence for a given voxel i as $X_{i,:} \in \mathbb{C}^{1\times L}$, using a Matlab style notation for indexing. Similarly, we will denote the response image at a given time t by the column vector $X_{:,t} \in \mathbb{C}^N$.

2.2. Estimating the Bloch parameters from the responses

In order to be able to retrieve $\{\rho_i, \theta_i\}$, it is necessary that the excitation sequence (α, TR) is "sufficiently rich" so that the voxel's dynamics $X_{i,:}$ is identifiable (random sequences (α, TR) seem to suffice in practice). Mathematically, this means that there is an embedding of $\mathbb{R}_+ \times \mathcal{M}$ into $\mathbb{C}^{1\times L}$.[2] We will call $\mathcal{B} = B(\mathcal{M}; \alpha, \mathrm{TR}) \subset \mathbb{C}^{1\times L}$ the Bloch response manifold and denote its positive cone by $\mathbb{R}_+\mathcal{B}$.

Inferring $\{\rho_i, \theta_i\}$ from the sequence $X_{i,:}$ can be achieved by locating $X_{i,:}$ on $\mathbb{R}_+\mathcal{B}$ and evaluating the associated parameters. This can be approximated in practice by projecting onto the cone of a discretized version of the response manifold.

Let $\theta_i^{(k)} = \{T_1^{(k)}, T_2^{(k)}, \delta f^{(k)}\}_{k=1:P}$ be a discrete sampling of the parameter space $\mathcal{M}$ to a desired accuracy. We then construct a "dictionary" $D \in \mathbb{C}^{P\times L}$ of the magnetization responses $D_k = B(\theta_i^{(k)}; \alpha, \mathrm{TR})$, $k = 1, \ldots, P$. We also construct a look-up table (LUT) to provide an approximate inverse for $B(\theta_i; \alpha, \mathrm{TR})$ such that $\theta_i^{(k)} = \mathrm{LUT}_B(k)$.

The approximate orthogonal projection onto the cone of response manifold $\mathbb{R}_+\mathcal{B}$, denoted by $\tilde{\mathcal{P}}_{\mathbb{R}_+\mathcal{B}}$, satisfies $\tilde{\mathcal{P}}_{\mathbb{R}_+\mathcal{B}}(X_{i,:}) = \hat{\rho}_i D_{\hat{k}_i}$, where

$$\hat{k}_i \in \operatorname*{argmax}_{1\le k\le P} \frac{\mathrm{Real}(\langle D_k, X_{i,:}\rangle)}{\|D_k\|_2}, \tag{2}$$

and $\hat{\rho}_i = \max\{\mathrm{Real}(\langle D_{\hat{k}_i}, X_{i,:}\rangle)/\|D_{\hat{k}_i}\|_2^2, 0\}$. The Bloch parameters corresponding to $X_{i,:}$ satisfies $\hat{\theta}_i = \mathrm{LUT}_B(\hat{k}_i)$.

3. MRF IMAGING

For the complete spatial image, we have $\theta \in \mathcal{M}^N$ and $\rho \in \mathbb{R}_+^N$. For convenience, let us denote the full mapping of the product space as $X = f(\rho, \theta)$, with $f : \mathbb{R}_+^N \times \mathcal{M}^N \to (\mathbb{R}_+\mathcal{B})^N \subset \mathbb{C}^{N\times L}$.

Unfortunately, it is impractical to observe the full spatial magnetization $X_{:,t}$ at each repetition time t within the necessary time window. It is necessary to resort to some form of undersampling. We can therefore define the observation sequence $Y_{:,t} \in \mathbb{C}^M$ as

$$Y_{:,t} = P(t) F X_{:,t}, \tag{3}$$

where $F \in \mathbb{C}^{N\times N}$ represents the forward discrete Fourier transform, and $P(t) \in \{0,1\}^{M\times N}$ is a t-dependent projection onto a subset of the output coefficients. We will denote the full linear observation mapping from the spatial magnetization sequence to observation sequence as $Y = h(X)$.

[2] Strictly speaking, this is only an embedding for $\rho_i > 0$.

3.1. MRF Matched filter reconstruction

In [1], the image sequence is reconstructed using back projection which is given by[3]

$$\hat{X}_{:,t} = F^H P(t)^T Y_{:,t}. \tag{4}$$

Due to the high level undersampling, this process generates extreme aliasing and therefore very noisy images. However, Ma *et al.* argue that by projecting each voxel sequence onto the Bloch response dictionary D, the noise can be suppressed and relatively clean parameter maps can be generated.

The procedure works through a form of noise averaging. Although each individual image is very noisy, the noise is greatly reduced when the voxel sequences are projected onto the Bloch response manifold. However, this ignores the main tenet of CS: aliasing is not noise but interference and under the right circumstances it can be completely removed. We explore this idea next.

3.2. Compressed Quantitative MRI

In order to be able to retrieve $\{\rho, \theta\}$ from Y, we propose a CS solution that has three key ingredients: a random pulse excitation sequence following the original MRF technique; a random subsampling strategy; and an efficient iterated projection algorithm [2] that imposes consistency with the Bloch equations.

3.2.1. Bloch response recovery via iterated projection

In [2], the Projection Landweber Algorithm (PLA) was proposed as an extension of the popular Iterated Hard Thresholding Algorithm [5, 6]. PLA is applicable to arbitrary union of subspace models as long as we have access to a computationally tractable projection operator within the complete signal space. In our case, the ideal algorithm is given by the recursion

$$X^{(n+1)} = \mathcal{P}_{(\mathbb{R}_+\mathcal{B})^N}\left[X^{(n)} + \mu\, h^H\Big(Y - h\Big(X^{(n)}\Big)\Big)\right], \tag{5}$$

where n is the recursion index, $\mathcal{P}_{(\mathbb{R}_+\mathcal{B})^N}$ is the orthogonal projection onto the signal model $(\mathbb{R}_+\mathcal{B})^N$, and μ is a stepsize.

In practice, we replace the projection $\mathcal{P}_{(\mathbb{R}_+\mathcal{B})^N}$ by the approximate orthogonal projection, denoted by $\tilde{\mathcal{P}}_{(\mathbb{R}_+\mathcal{B})^N}$, computed by separately projecting the individual voxel sequences $X_{i,:}^{(n)}$ using the projector $\tilde{\mathcal{P}}_{\mathbb{R}_+\mathcal{B}}$ defined in Section 2.2. We call the resulting algorithm BLIP (BLoch response recovery via Iterated Projection).

The current theory for PLA guarantees stable recovery as long as h satisfies a so-called Restricted Isometry Property

[3] This is actually only an approximation since in [1] the authors use a nonuniform Fourier transform since their spiral read out does not lie on the DFT grid.

(RIP) for the signal model. That is, if there exists a constant $\delta > 0$ such that

$$(1-\delta)\|X-\tilde{X}\|_2^2 \leq \tfrac{N}{M}\|h(X-\tilde{X})\|_2^2 \leq (1+\delta)\|X-\tilde{X}\|_2^2 \quad (6)$$

for all pairs X and $\tilde{X}$ in $(\mathbb{R}_+\mathcal{B})^N$. We will describe how to achieve such an embedding later in section 3.2.2. Note that the vectors of interest in the above RIP are the *chords* of $\mathbb{R}_+\mathcal{B}$, *i.e.*, the vectors belonging to $U := \{\mathbb{R}_+\mathcal{B} - \mathbb{R}_+\mathcal{B}\}\backslash\{0\}$.

The theory [2] also requires that $M(1+\delta)/N < 1/\mu < 3M(1-\delta)/(2N)$ for the successful recovery. If h is essentially "optimal", *e.g.*, a random ortho-projector, then we should set the stepsize $\mu \approx N/M$ since in the large system limit $\delta \to 0$. In practice, selection of the correct stepsize is crucial in order to attain good performance from these iterative projection based algorithms [7, 2]. While the matched filter used in [1] can be interpreted as a single iteration of PLA with $\mu = 1$, a substantially more aggressive step size is proposed by the theory and results in significant improvements. In practice, we noticed that it is also beneficial to select adaptively the step size for PLA to ensure stability. In the experiments, we use an adaptive stepsize selection as in [7].

3.2.2. Strategies for subsampling k-space

In order to ensure guaranteed parameter map recovery, recall that (1) the excitation response mapping f must be an embedding (achieved using a random excitation sequence (α, TR)), and (2) the sampling operator h should satisfy a suitable RIP. We will see that this latter condition is satisfied using an appropriate undersampling strategy, if the vectors in U satisfies certain conditions.

Since we only take a small number of measurements at each repetition time, we cannot expect to achieve a stable embedding without imposing further constraints on the excitation response. For example, if the embedding f was induced in the first few repetition times and all further responses were non-informative, we would not have taken sufficient measurements from the informative portion of the response. Therefore, we need to consider responses that somehow spread the information across the repetition times. We quantify the spread of the information for the excitation response through the *flatness* of the *chords* of $\mathbb{R}_+\mathcal{B}$.

Definition 1 *Let U be a collection of vectors $\{u\}$ in $\mathbb{C}^L$. We denote the* flatness *of the these vectors by*

$$\lambda := \max_{u \in U} \|u\|_\infty / \|u\|_2. \quad (7)$$

Note that from standard norm inequalities $L^{-1/2} \leq \lambda \leq 1$.

The chords $u \in U$ of $\mathbb{R}_+\mathcal{B}$ are considered sufficiently flat if $\lambda \sim L^{-1/2}$. In Fig. 1, we present a plot λ^{-2}/L as a function of sequence length L for an example of the response differences. From this plot, it can be deduced that λ^{-2} grows roughly proportionally to L.

In constructing our measurement function h, we note that the signal model contains no spatial structure. Therefore, we should expect to uniformly sample k-space in order to achieve an embedding, in contrast to the variable density sampling of [1] (experiments in [9] support this argument). As no spatial structure is considered, we describe, for simplicity, our sampling strategy for a 1D signal of N pixels. Let $k_1, \ldots, k_N \in \{0, \ldots, N-1\}$ denote the N measurable discrete frequencies. Then, we define a random measurement operator by $P(\zeta_t)F$, where $(\zeta_t)_{1 \leq t \leq L}$ is a sequence of independent random variables uniformly drawn from $\{0, \ldots, p-1\}$ and $P(\zeta_t) \in \{0,1\}^{M \times N}$ has entries

$$(P(\zeta_t))_{i,j} = \begin{cases} 1 & \text{if } k_j = (i-1)p + \zeta_t, \\ 0 & \text{otherwise.} \end{cases}, \quad (8)$$

with $i = 1, \ldots, M$ and $j = 1, \ldots, N$. For convenience, we assume that p divides N exactly so that $N = pM$. In words, we regularly subsample the k-space by a factor of p with random shifts of the selected samples across time. For 2D images, this corresponds to a regular subsampling of the k-space in one direction and a complete sampling of the selected lines in the other direction. This can be achieved using a randomized version of multishot *Echo-Planar Imaging* (EPI) [8].

The following theorem shows that random EPI along with an excitation response with appropriate chord flatness is sufficient to provide us with a measurement operator, h, that satisfies the RIP on our signal model (proof available in [9]).

Theorem 1 (RIP for random EPI) *Given an excitation response cone $\mathbb{R}_+\mathcal{B}$ of dimension $d_\mathcal{B}$, whose chords have a flatness λ, and a random EPI operator $h : (\mathbb{R}_+\mathcal{B})^N \to \mathbb{C}^{M \times L}$. With probability at least $1-\eta$, h is a restricted isometry on $(\mathbb{R}_+\mathcal{B})^N - (\mathbb{R}_+\mathcal{B})^N$ with constant δ as long as*

$$\lambda^{-2} \geq C\delta^{-2}p^2 d_\mathcal{B} \log(N/\delta\eta), \quad (9)$$

for some constant C independent of $p, N, d_\mathcal{B}, \delta$ and η.

Specifically, if $\lambda \sim L^{-1/2}$ then we require $L \sim p^2 d_\mathcal{B}$ excitation pulses. While we might hope to get L of the order of $pd_\mathcal{B}$ it appears that this is not possible, at least for a worst case RIP analysis based on the flatness criterion alone. Indeed, in the experimental section, we will provide evidence to suggest that $L \sim p^2$ is indeed the scaling behavior that we empirically observe.

4. EXPERIMENTS

4.1. Setting

In order to test the efficacy of our method, we performed a set of simulations using an anatomical brain phantom adapted from the anatomical brain phantom of [10], available at the BrainWeb repository.[4] The image contains 256×256 pixels

[4] http://brainweb.bic.mni.mcgill.ca/brainweb/

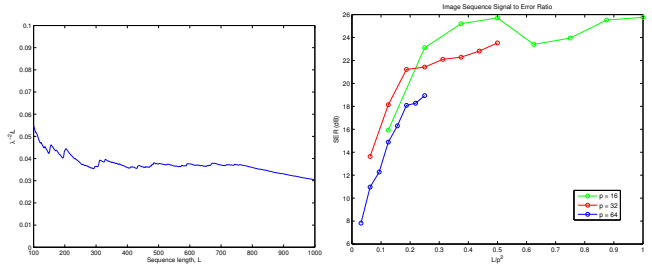

Fig. 1. Left: λ^{-2}/L as a function of L for some example of response differences. Right: plot of the image sequence SER (dB) against L/p^2 for three different levels of undersampling $p = 16$ (green), $p = 32$ (red), and $p = 64$ (blue).

and was restricted to 6 material components. The material properties were chosen to be representative of the correct tissue type [11] and were set so that there is not an exact match to the sampling of the Bloch response manifold. The proton densities were fixed to give little discrimination for individual parameters.

For the excitation sequences, we use IR-SSFP sequences [1] with random flip angles drawn from an i.i.d. Gaussian distribution with standard deviation of 10 degrees. The repetition times were uniformly spaced at an interval of 10 ms. The Bloch response manifold was sampled in a similar manner to [1], however, we have only considered variation in T1 and T2 here, assuming that δf is equal to zero. This results in a dictionary of size $3379 \times L$, its range spanning the values for the expected tissue types. For the Fourier subsampling, we use the random EPI sampling scheme detailed in Section 3.2.2.

4.2. Results

To get a visual indication of the performance of the BLIP approach over the original MRF reconstruction, 3 different parameter estimates for $L = 200$ and $p = 16$ are given in Fig. 2. The left hand column shows the ground truth parameter maps while the middle row shows the MRF reconstruction (1 iteration of the BLIP algorithm with $\mu = N/M$) and the right hand column shows the BLIP estimates. While the main aspects of the parameter maps are visible in the MRF reconstructions, there are still substantial aliasing artifacts. These are most prominent in the density and T_1 estimates. In contrast, the BLIP estimates are virtually distortion-free, indicating that good spatial parameter estimates can be obtained with as little as 200 excitation pulses.

In the second experiment, we evaluate the image sequence signal-to-error-ratio[5] (SER) as a function of L and p. Recall that the theory suggested that this performance might degrade roughly as a function of L/p^2. Fig. 1 shows a plot of the image sequence SER as a function of L/p^2 for three different subsampling rates. We can see that the rapid growth of the SER that we associate with successful recovery occurs in each case at roughly the same value of L/p^2. This seems to suggest that the predicted scaling behaviour for L and p in random EPI to achieve RIP is of the right order.

[5]This is calculated as $20 \log_{10}(\|X - \hat{X}\|_F / \|X\|_F)$ for a target signal X with the estimate $\hat{X}$.

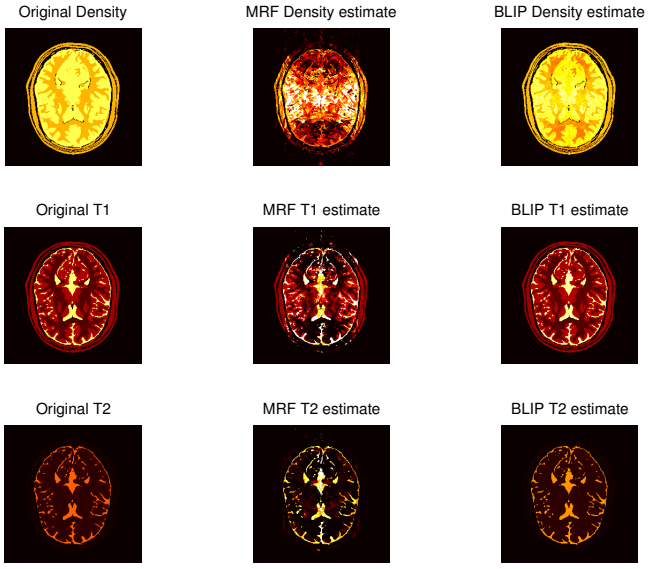

Fig. 2. The top row shows the density maps: original (left), MRF estimate (center) and BLIP estimate (right). The middle and bottom rows shows the T1 and T2 maps respectively in the same order.

5. CONCLUSIONS

We have presented a principled mathematical framework for compressed quantitative MRI based around the recently proposed technique of MRF [1]. The key elements of our approach have been: the characterization of the signal model through the Bloch response manifold; the identification of a provably good reconstruction algorithm based on iterative projection; an excitation response condition based on a newly introduced measure of *flatness*; and a random EPI scheme that has the necessary RIP condition.

While the current work is targeted at a CS framework for MRF, we believe that many elements of it should be more broadly applicable. Specifically, the RIP condition for randomized EPI may have applications in other MR imaging strategies, and the characterization of excitation response in terms of flatness could be a useful tool for the analysis of other CS schemes involving some form of active sensing.

6. REFERENCES

[1] D. Ma, V. Gulani, N. Seiberlich, K. Liu, J. L. Sunshine, J. L. Duerk, and M. A. Griswold, "Magnetic resonance fingerprinting," *Nature*, vol. 145, pp. 187–192, 2013.

[2] T. Blumensath, "Sampling and reconstructing signals from a union of linear subspaces," *IEEE Trans. Inf. Theory*, vol. 57, no. 7, pp. 4660 – 4671, 2001.

[3] B.A. Hargreaves, S.S. Vasanawala, J.M. Pauly, and D.G. Nishmura, "Characterization and reduction of the transient response in steady-state MR imaging," *Magn. Reson. Med.*, vol. 46, no. 1, pp. 149–158, 2001.

[4] C. Ganter, "Off-resonance effects in the transient response of SSFP sequences," *Magn. Reson. Med.*, vol. 52, no. 2, pp. 368–375, 2004.

[5] T. Blumensath and M. E. Davies, "Iterative hard thresholding for sparse approximation," *J. Fourier Anal. Appl.*, vol. 14, no. 5-6, pp. 629–654, 2008.

[6] T. Blumensath and M. E. Davies, "Iterative hard thresholding for compressed sensing," *Appl. Comput. Harmon. Anal.*, vol. 27, no. 3, pp. 265–274, 2009.

[7] T. Blumensath and M. E. Davies, "Normalised iterative hard thresholding; guaranteed stability and performance," *IEEE J. Sel. Top. Signal Process.*, vol. 4, no. 2, pp. 298–309, 2010.

[8] G.C. McKinnon, "Ultrafast interleaved gradient-echo-planar imaging on a standard scanner," *Magn. Reson. Med.*, vol. 30, no. 5, pp. 609–616, 1993.

[9] M. E. Davies, G. Puy, P. Vandergheynst, and Y. Wiaux, "A compressed sensing framework for magnetic resonance fingerprinting," *eprint* arXiv:1312.2465, 2013.

[10] D.L. Collins, A.P. Zijdenbos, V. Kollokian, J.G. Sled, N.J. Kabani, C.J. Holmes, and A.C. Evans, "Design and construction of a realistic digital brain phantom," *IEEE Trans. Med. Imaging*, vol. 17, no. 3, pp. 463–468, 1998.

[11] J.P. Hornak, "The basics of MRI," Available online at: `http://www.cis.rit.edu/htbooks/mri/`.