Image Encoding and Reconstruction with Unconventional MRI Hardware

Daniel Abraham

Abstract

Magnetic Resonance Imaging (MRI) is traditionally performed in a confined bore with a strong permanent magnetic field. Since these systems are expensive and have reduced accessibility, there has been a push to use lower field strengths, which directly lowers the cost of the scanners. Due to the inherently low SNR of low field systems, excessive temporal averaging (i.e. long scan times) are necessary to improve the SNR. This is uncomfortable since the subject must stay stationary to avoid motion artifacts. We investigate the use of wearable MRI for image encoding at low field strengths. The wearable hardware combined with imaging during the 6-8 hour human sleep cycle will provide the SNR benefit of temporal averaging in a comfortable way, without sensitivity to motion. Such hardware results in a non-trivial encoding model. In this work we will derive the new encoding model for the simple proposed wearable hardware system. We use simple design choices for our wearable hardware and evaluate the regularized reconstructions at moderate levels of noise.

1 Introduction

Magnetic Resonance Imaging (MRI) is typically performed in the presence of an extremely strong and uniform magnetic field. Such fields are necessary to strongly polarize spins (tiny magnetic fields produced by hydrogen atoms) to produce the sufficient SNR required to image within reasonable time constraints (less than an hour). Image encoding is performed with spatially linear varying magnetic fields produced by gradient coils. By modulating the strength of these gradient coils over the three coordinate axes, one can acquire all necessary points in the spatial frequency domain (k-space), and reconstruct using an Inverse Fourier Transform. In addition to linear spatial encoding, multiple MRI RF receiver coils are often placed throughout the surface of the subject to create additional image encoding with coil-sensitivity based spatial weighting maps.

Having powerful magnetic fields (1 to 7 Tesla) are great for achieving a high SNR however they are the main reason why MRIs are expensive, uncomfortable, and have a reduced accessibility globally. In order to address these pitfalls, there have been many proposed low field systems (less than 1 Tesla) that operate at a reduced field strength [1], [2]. Such systems instead use opt for cheaper permanent magnets to generate the strong main field, and keep all other encoding hardware relatively unchanged. While this does indeed reduce the cost, the primary challenge is the reduced SNR at lower fields.

The most straightforward way to gain the lost SNR is to simply image for a longer time (i.e. averaging). However, most subjects are unable to sit still for such a long period of time (more than 30 minutes). Thus, many have applied deep learning techniques to low field images in order to denoise and improve image quality to that of high field systems. While this does help, the use of deep learning for denoising can only go so far before hallucinations become apparent – which is very dangerous in medical imaging applications. Motivated by the fundamental SNR wall at low field, we turn to a slightly different strategy:

Can we move all the MRI hardware onto the subject as a wearable, and acquire data during sleeping hours?

Assuming that such a system did exist, and was indeed comfortable enough to wear, then the SNR lost due to using a low field system can be bought back by imaging over the course of the human sleep cycle (6-8 hours). Such systems have not yet been built, but can be simulated with some fairly simple wearable MRI coil designs. Placing the imaging coils on the subject's head will result in spatially non-linear phase encoding, in contrast to spatially linear phase encoding (i.e. sampling in the spatial frequency domain). **This presents an opportunity to create a non-Fourier based image reconstruction algorithm for MRI with unconventional hardware, and pair it with strong image priors to help with denoising.** A visual description of the wearable setup is provided in Figure 1.

2 Related Work

Unconventional encoding techniques typically show up in the literature to either accelerate existing hardware at high field, or to enable lower cost hardware. We list a few of these ideas below:

- Pre-polarized MRI [3] is a method in which a strong transient magnetic field is used only to generate spin polarization, which is needed in order to receive strong signals, and is then switched off before encoding fields are used for data

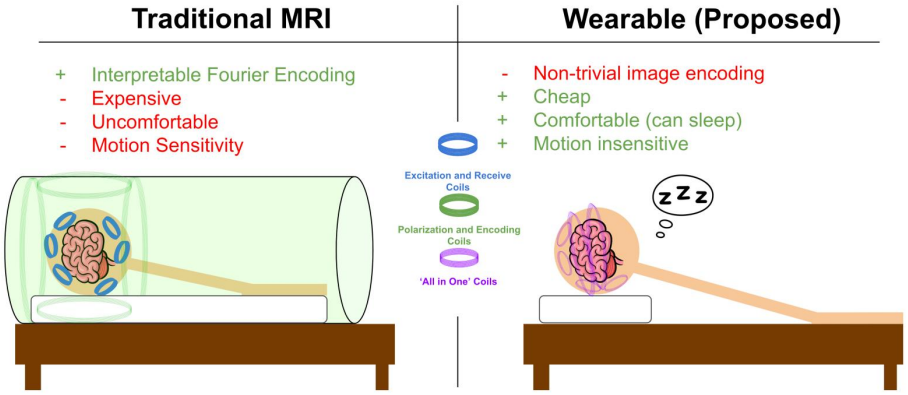

Figure 1: Comparison of a conventional MRI system (left) with the proposed wearable system (right).

acquisition. This is different than traditional MRI, where the polarizing main magnetic field is always on. In this work, we will use such a pre-polarizing field to achieve high spin polarization, and only augment the encoding coils.

- The authors of [1] construct a low field where the main magnetic field has a fixed pattern using a halbach array of permanent magnets, and this field is *mechanically* rotated in order to produce sufficient spatial encoding. Our idea is heavily inspired from this – we are instead opting for an *electrically* rotating magnetic field that is instead comfortably worn on the subject to justify scanning overnight.
- The authors of [4] replace conventional MRI gradient coils (which produce linearly varying fields along coordinate axes) with a set of coils that produce a non-linear field variation over space. They place the coils around the head of a subject, giving excellent spatial encoding around the extremities of the head, but very poor encoding in the center. In this work, a large uniform main field is used for spin polarization. In our case, we would be extending the encoding coils to also provide the main magnetic field needed for polarization, enabling low field applications.
- The authors of [5] allow for a variable main magnetic field, which is controlled by changing the applied current to such a field. The authors decided to target for spatially linearly varying magnetic fields, and perform additional image encoding using an additional set of RF coils. We extent this by exploring wearable non-linearly varying magnetic fields.

3 Theory and Methods

3.1 MRI Physics

In order to receive strong MR signals, a strong magnetic field is required to achieve spin polarization. A strong field aligns a number of hydrogen nuclei with the field direction in proportion to the field strength. Thus, in order to use more hydrogen nuclei, a strong field is necessary. In this work, we achieve the strong polarizing field using pre-polarized MRI [3]. Unlike conventional MRI systems, pre-polarized systems use transient polarizing fields, that are only active before and after encoding/data acquisition periods. This is achieved by pulsing large loop coils. We will use a set of two coils, mounted on the subject's head, to achieve a spin pre-polarization $\mathbf{M}_0(\mathbf{r})$, which describes the polarization induced by the pre-polarized coils as a function of the spatial coordinates $\mathbf{r} \in \mathbb{R}^3$.

In our imaging setup, we will have a set of K coils placed on the subjects head at several known locations, orientations, and diameters. Each coil is used both for generating encoding magnetic fields and receiving data. If one were to apply currents $I_1, \cdots, I_K$ to each coil, the magnetic field induced by all of these coils $\mathbf{B}(\mathbf{r}) : \mathbb{R}^3 \to \mathbb{R}^3$ is simply a linear combination of the unit field response $\mathbf{B}_k(\mathbf{r}) : \mathbb{R}^3 \to \mathbb{R}^3$ per coil

$$\mathbf{B}(\mathbf{r}) = \sum_k I_k \mathbf{B}_k(\mathbf{r}). \tag{1}$$

The encoding field $\mathbf{B}(\mathbf{r})$ will cause the polarized magnetic moment vectors $\mathbf{M}_0(\mathbf{r})$ to rotate about the field vector $\mathbf{B}(\mathbf{r})$ with a rotational frequency given by $\gamma||\mathbf{B}(\mathbf{r})||_2$, where γ is a constant known as the gyromagnetic ratio. As this magnetic moment rotates, it produces an RF signal which can be received by each of the K coils. Critically, the received signal is not at DC, and hence the coils can be used both to transmit DC magnetic fields, as well as receive AC signals in tandem.

Data is usually acquired after some prescribed echo time in order to allow tissue parameters to induce a contrast weighted magnetization $\tilde{\mathbf{M}}(\mathbf{r}) = f(\mathbf{M}_0(\mathbf{r}), T_1(\mathbf{r}), T_2(\mathbf{r}), \cdots)$, where $T_1(\mathbf{r}), T_2(\mathbf{r})$ are spatially varying maps of tissue relaxation parameters. There are many more medically relevant tissue parameters that will induce a contrast on the image which we omit,

but are well described CITE HERE. More importantly, in order to get medically useful images, our goal is to estimate $\tilde{\mathbf{M}}(\mathbf{r})$. Using some MRI physics and Faraday's law of induction, the signal received by the k^{th} coil is then given by

$$s_k(t) = -\frac{\partial}{\partial t}\int_{\mathbf{r}} \mathbf{B}_k(\mathbf{r}) \cdot \left(\mathbf{R}_{\hat{\mathbf{B}}}\big(\gamma||\mathbf{B}(\mathbf{r})||_2 t\big)\tilde{\mathbf{M}}(\mathbf{r})\right) d\mathbf{r}, \tag{2}$$

where $\mathbf{R}_{\mathbf{v}}(\theta)$ describes a rotation operator about $\mathbf{v}$ by angle θ. **Importantly, the received signal is in fact linear in $\tilde{\mathbf{M}}(\mathbf{r})$.**

3.2 Linear System Formulation

Upon vectorizing the multi-coil received temporal signals $s_k(t) \to \mathbf{b} \in \mathbb{C}^M$ and the desired contrast weighted magnetization vectors $\tilde{\mathbf{M}}(\mathbf{r}) \to \mathbf{m} \in \mathbb{C}^N$, we can combine (1) and (2) to arrive at a much simpler linear system, where the forward model depends on the applied currents $I_1, \cdots, I_K$:

$$\mathbf{A}(I_1, \cdots, I_K)\mathbf{m} + \mathbf{n} = \mathbf{b} \tag{3}$$

where $\mathbf{A} \in \mathbb{C}^{M\times N}$ is the discretized forward model described in (2), and $\mathbf{n} \sim \mathcal{N}(0, \sigma^2\mathbf{I})$ is Gaussian noise induced by resistance in the coils and subject. Since we are analyzing the analytic signal in each coil via IQ demodulation, the received data is complex. A visual representation of $\mathbf{A}$ for some fixed set of currents is shown in Figure 2.

In order to have sufficient encoding to invert (3), and battle the relatively large levels of Gaussian noise at low field, we can repeat our imaging experiment with a set of N currents $\{(I_1^{(i)}, \cdots, I_K^{(i)})\}_{i=1}^N$, and attempt to solve the following regularized inverse problem:

$$\min_{\mathbf{m}} || \underbrace{\begin{bmatrix} \mathbf{A}(I_1^{(1)}, \cdots, I_K^{(1)}) \\ \vdots \\ \mathbf{A}(I_1^{(N)}, \cdots, I_K^{(N)}) \end{bmatrix}}_{\mathbf{A}} \mathbf{m} - \underbrace{\begin{bmatrix} \mathbf{b}^{(1)} \\ \vdots \\ \mathbf{b}^{(N)} \end{bmatrix}}_{\mathbf{b}} ||_2^2 + \mathcal{R}(\mathbf{m}). \tag{4}$$

When implementing the algorithms above, we use an MRI experiment setup with two pre-polarizing coils, $K = 8$ encoding/receive coils, a 2D brain phantom with an effective diameter of 22cm, and maximum current limits of 10 amps for both pre-polarizing and encoding coils. The biot-savart law is used to simulate the coil fields for all polarizing and encoding coils. Pytorch was used to implement the discussed linear operators and iterative algorithms. We use $N = 200$ measurements, where each measurement consists of a data acquisition window with 2ms at a sampling rate of $\Delta t = 4\mu$s. The imaging currents were randomly chosen from a uniform distribution $I_k^{(i)} \sim \mathcal{U}(-I_{\text{max}}, I_{\text{max}})$.

4 Analysis and Evaluation

In this work, we will perform four reconstructions to analyze and evaluate our wearable low-field encoding system:

- **CG**: Conjugate Gradient reconstruction with tikonov hyper-parameter λ_2:

$$\min_{\mathbf{m}} ||\mathbf{A}\mathbf{m} - \mathbf{b}||_2^2 + \lambda_2||\mathbf{m}||_2^2$$

- **TV**: Total Variation regularized ADMM reconstructions with ℓ_1 hyper-parameter λ_1:

$$\begin{aligned} \mathbf{m} &\leftarrow \operatorname*{argmin}_{\mathbf{m}} \frac{1}{2}||\mathbf{A}\mathbf{m} - \mathbf{b}||_2^2 + \frac{\rho}{2}||\mathbf{T}\mathbf{V}\mathbf{m} - \mathbf{z} + \mathbf{u}||_2^2 \\ \mathbf{z} &\leftarrow \mathcal{S}_{\lambda_1/\rho}(\mathbf{T}\mathbf{V}\mathbf{m} + \mathbf{u}) \\ \mathbf{u} &\leftarrow \mathbf{u} + \mathbf{T}\mathbf{V}\mathbf{m} - \mathbf{z} \end{aligned}$$

- **WAV**: Wavelet regularized ADMM reconstructions with ℓ_1 hyper-parameter λ_1:

$$\begin{aligned} \mathbf{m} &\leftarrow \operatorname*{argmin}_{\mathbf{m}} \frac{1}{2}||\mathbf{A}\mathbf{m} - \mathbf{b}||_2^2 + \frac{\rho}{2}||\mathbf{W}\mathbf{m} - \mathbf{z} + \mathbf{u}||_2^2 \\ \mathbf{z} &\leftarrow \mathcal{S}_{\lambda_1/\rho}(\mathbf{W}\mathbf{m} + \mathbf{u}) \\ \mathbf{u} &\leftarrow \mathbf{u} + \mathbf{W}\mathbf{m} - \mathbf{z} \end{aligned}$$

- **DnCNN**: Denoise CNN regularized ADMM reconstructions with hyper-parameter ρ:

$$\begin{aligned} \mathbf{m} &\leftarrow \operatorname*{argmin}_{\mathbf{m}} \frac{1}{2}||\mathbf{A}\mathbf{m} - \mathbf{b}||_2^2 + \frac{\rho}{2}||\mathbf{m} - \mathbf{z} + \mathbf{u}||_2^2 \\ \mathbf{z} &\leftarrow \mathcal{D}(\mathbf{m} + \mathbf{u}) \\ \mathbf{u} &\leftarrow \mathbf{u} + \mathbf{m} - \mathbf{z} \end{aligned}$$

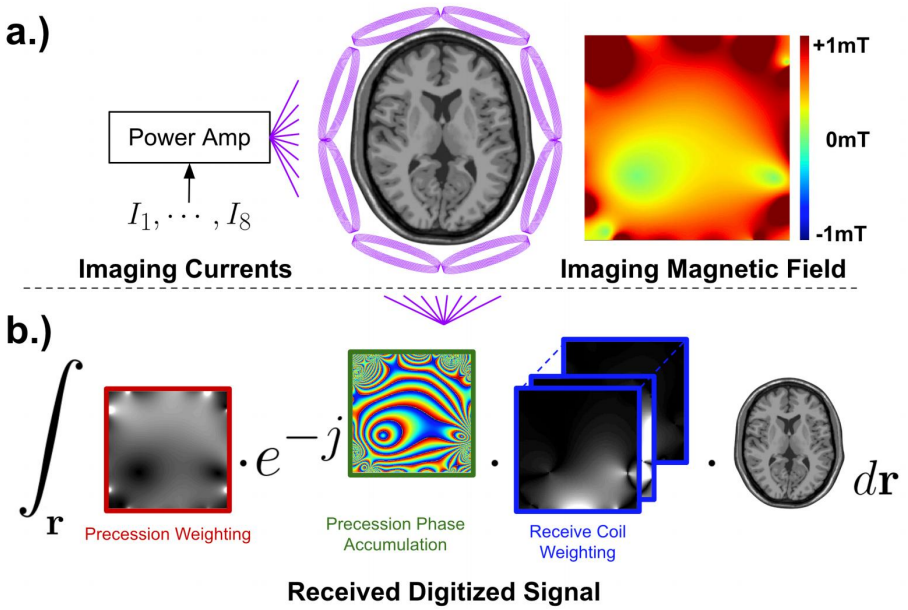

Figure 2: a.) Currents $I_1, \cdots, I_8$ are applied to their respective coils to generate an imaging magnetic field (right). b.) The digitized signal induced on each coil is given by a linear function of the image. The encoding functions are dependent on the imaging magnetic field, and hence they are also dependent on the currents $I_1, \cdots, I_8$. The same set of $K = 8$ coils are used for encoding and receiving.

The Conjugate Gradient [6] algorithm is used for CG and for the $\mathbf{m}$ update of TV, Wavelet, and DnCNN. λ_2 is the tikonov hyper-parameter controlling the ℓ_2 regularization for CG. We use the ADMM [7] formulation to solve TV, WAV, and DnCNN methods, as shown in the update rules. $\mathbf{TV}$ is a finite difference operator, $\mathbf{W}$ is a wavelet transform, and $\mathcal{D}$ is the DnCNN described in [8]. The ADMM algorithm is described in [9].

5 Results

All reconstructions were run on an NVIDIA A5000 GPU, equip with 24GB of GPU RAM. Performing ADMM reconstructions took roughly 15 minutes, while CG reconstructions took 2 minutes. Figure 3 demonstrates the un-regularized CG reconstruction without any noise added to the system. 500 iterations of the CG algorithm were used to reconstruct the image. Regularized reconstructions are then run in the presence of moderate Gaussian noise. We fixed the noise level to $\sigma = 0.002$, swept over the relevant hyper-parameter for each case, and compare the resulting reconstruction to the ground truth brain phantom. These parameter sweeps are shown in Figure 4. From this figure, we select the hyper-parameters, and fine tune them for best visual results. The final hyper-parameters are used in Figure 5, where we compare all the methods for the same noise level.

6 Discussion and Conclusion

The previous figures show that image recovery with a moderate amount of Gaussian noise is feasible for our wearable MRI system under ideal conditions. Various regularization functions, given a proper choice of hyper-parameters, help reduce the impact of noise on the final images. One of the major challenges in this work was managing computation and choosing the reconstruction hyper-parameters. The forward model $\mathbf{A}$ is quite expensive to compute, resulting in very long reconstruction times. As a result, not all hyper-parameters are truly optimally chosen. For example, the ADMM algorithm has two hyper-parameters, however in our examples we only optimized one of them. Thus, improving image reconstruction times is an important future goal, which will aid in choosing hyper-parameters and general ease of use.

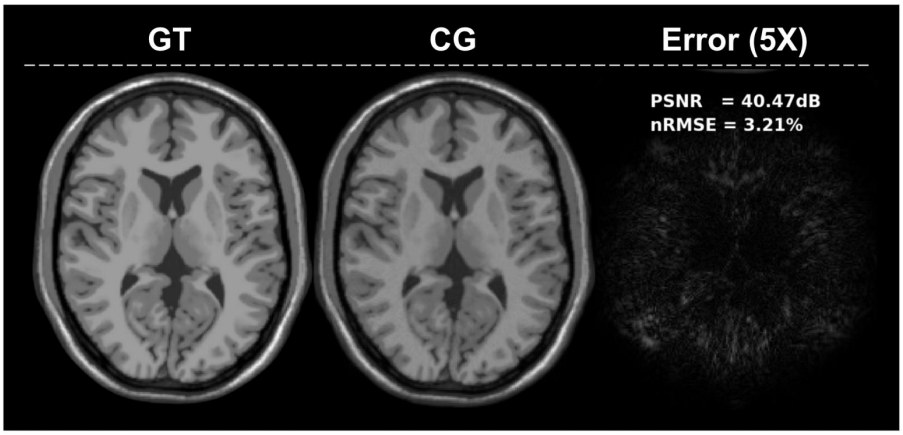

Figure 3: Noiseless unregularized reconstruction. Ground Truth (left), Conjugate Gradient method (middle) and 5X amplified difference image (right), with PSNR and nRMSE of 40.47dB and 3.21% respectively.

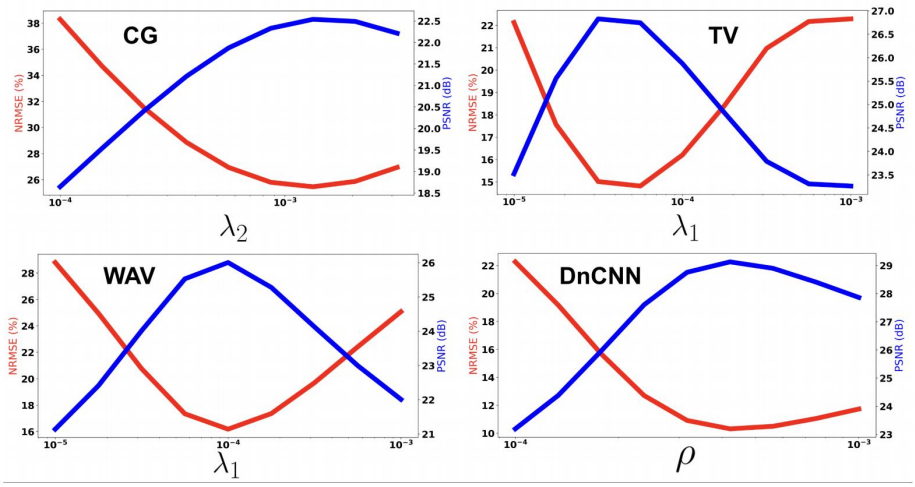

Figure 4: Hyper parameter sweeps are shown with Conjugate Gradient ℓ_2 regularization (top left), Total Variation ℓ_1 regularization (top right), Wavelet ℓ_1 regularization, and DnCNN.

From the noiseless case, we can see that image encoding is not an issue with such a wearable system. Imaging over 6-8 hours of sleep should provided tons of time to encode the image. The more difficult part is if such a system can produce enough SNR. The SNR that the wearable system is capable of producing at depends on the power applied to the coils and the coil size/geometry. While the increasing power and coil size improves SNR, it unfortunately reduces patient comfort due to weight and heating. Thus, it is important to explore the trade-off between comfortable hardware and sufficient SNR. In future work, we plan to characterize the precise noise characteristics of our wearable hardware system in order to identify exactly how long of a scan is needed in order to achieve reasonable SNR values without causing too much discomfort.

References

[1] Lawrence L Wald, Patrick C McDaniel, Thomas Witzel, Jason P Stockmann, and Clarissa Zimmerman Cooley. Low-cost and portable mri. *Journal of Magnetic Resonance Imaging*, 52(3):686–696, 2020. 1, 2

[2] W Taylor Kimberly, Annabel J Sorby-Adams, Andrew G Webb, Ed X Wu, Rachel Beekman, Ritvij Bowry, Steven J Schiff, Adam de Havenon, Francis X Shen, Gordon Sze, et al. Brain imaging with portable low-field mri. *Nature Reviews Bioengineering*, 1(9):617–630, 2023. 1

[3] Greig Scott, Blaine Chronik, N Matter, H Xu, P Morgan, L Wong, A Macovski, and S Conolly. A prepolarized mri scanner. In *Proc. Intl. Soc. Mag. Reson. Med*, volume 9, page 610, 2001. 1, 2

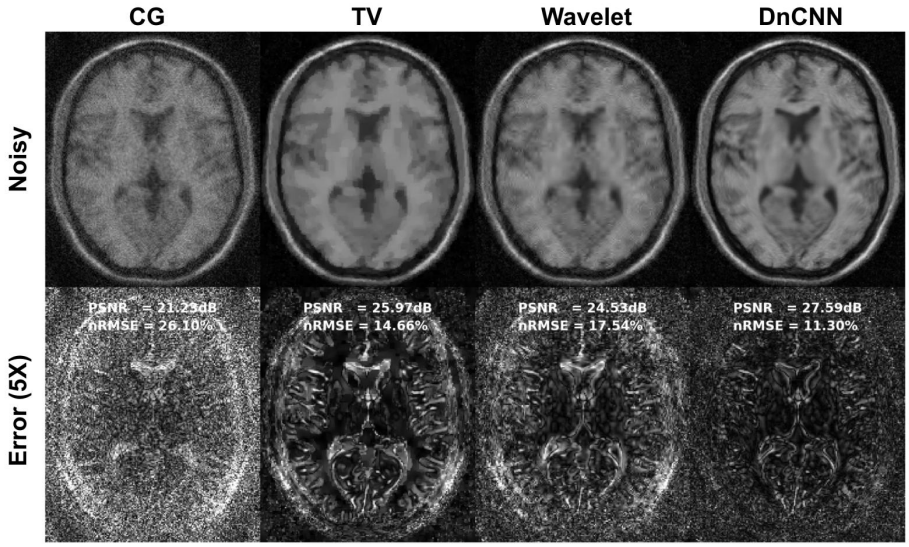

Figure 5: Noisy reconstructions shown for Congugate Gradient, Total Variation, Wavelet, and DnCNN (top) along with 5X amplified error images (bottom).

[4] Juergen Hennig, Anna Masako Welz, Gerrit Schultz, Jan Korvink, Zhenyu Liu, Oliver Speck, and Maxim Zaitsev. Parallel imaging in non-bijective, curvilinear magnetic field gradients: a concept study. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 21:5–14, 2008. 2

[5] Kartiga Selvaganesan, Yuqing Wan, Yonghyun Ha, Baosong Wu, Kasey Hancock, Gigi Galiana, and R Todd Constable. Magnetic resonance imaging using a nonuniform bo (nubo) field-cycling magnet. *Plos one*, 18(6):e0287344, 2023. 2

[6] Jonathan Richard Shewchuk et al. An introduction to the conjugate gradient method without the agonizing pain. 1994. 4

[7] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122, 2011. 4

[8] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017. 4

[9] Gordon Wetzstein. Stanford electrical engineering 367 computational imaging: Note 11. `https://stanford.edu/class/ee367/reading/ee367_notes_inverse_problems.pdf`. 4