

Deconvolution of Complex-Valued Partial-Fourier Images

Alexander R Toews
Department of Electrical Engineering
Stanford University
artoews@stanford.edu

Abstract

“Partial-Fourier” is an under-sampling technique employed in magnetic resonance imaging to reduce image acquisition time. The drawback of Partial-Fourier sampling is that it produces images which are blurred along one dimension. We formulate the problem of reconstructing Partial-Fourier images as a non-blind deconvolution problem and solve it using convex optimization. Total variation and Hessian-based priors are investigated. Image accuracy is improved by 3.5 and 4.5 dB over the default method of zero-filling missing data for two magnetic resonance images.

1. Introduction

Magnetic resonance imaging (MRI) is a safe, non-invasive imaging modality with excellent soft tissue contrast. MRI is routinely used in clinical settings to obtain cross-sectional images of live human subjects.

MRI differs from natural imaging in two key respects which are fundamental to the work that will be presented here.

First, the anatomical image of interest is complex-valued. That is to say that each image pixel has both a magnitude and phase ascribed to it. One powerful feature of MRI is that it can be dynamically programmed to encode information into each pixel’s magnitude and phase. This is why the same MRI machine can be used to investigate arthritis, neural activity, and cardiovascular flow without changing a single piece of hardware.

The second key difference between MRI and natural imaging is that raw MRI data is collected serially in the spatial frequency domain (referred to as k-space in MRI literature), rather than simultaneously in the primal image domain like a traditional camera. The anatomical image of interest is reconstructed by taking the inverse Fourier transform of the raw k-space data.

A significant challenge in MRI is that it takes a long time to (Nyquist) fully sample k-space at the spatial resolution needed for clinical use. In other words, sampling one point

of k-space at a time is relatively slow. It is common practice to undersample k-space to mitigate this challenge. The problem then becomes how best to reconstruct this under-sampled data.

In this work I focus on one such k-space undersampling scheme which is widely used in clinical practice: partial-Fourier (PF) sampling. In PF sampling, the k-space is fully sampled along all but one dimension. For this one dimension, frequencies above some threshold are only sampled on one side of k-space. An example PF sampling scheme is depicted in Figure 1.

If the missing frequency data is filled in with zeros, then an inverse Fourier transform produces an image which appears blurred along the PF dimension (e.g. Figure 2). From a computational imaging perspective, the PF sampling mask is an optical transfer function (OTF). Thus our image is in fact the result of convolving the fully sampled image with a complex-valued point spread function (PSF). The PSF here is a phase-modulated sinc function of the PF-sampled dimension (Figure 3).

2. Related Work

PF sampling has been around for a few decades now, and many schemes have been developed for reconstructing such data [7]. Standard methods for PF reconstruction [8] [3] try to exploit Hermitian symmetry in k-space. But these methods are very limited in scope because they require the image phase to be slowly varying, which is often not the case. More recently, a few papers have proposed using convolutional neural networks (CNNs) to learn the mapping from undersampled to fully sampled images in MRI ([5], [11], [10]), but none of these investigate the PF sampling scheme. The key challenge in this work is to investigate PF reconstruction for images for which current methods fall short - images which do not have slowly varying phase.

3. Methods

I choose to frame the PF reconstruction problem as a non-blind deconvolution. I proceed to find solutions to this

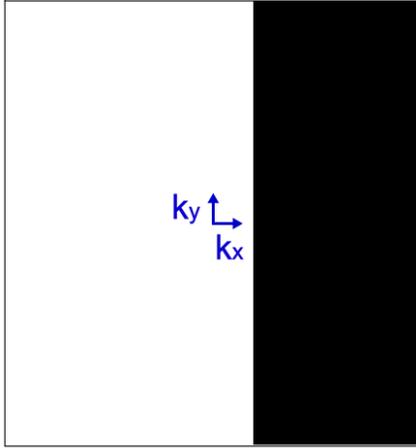


Figure 1. Example of a PF sampling scheme, where the sampled region is colored white. k_x and k_y represent the Cartesian basis for spatial frequencies in 2D. The PF sampling dimension is along x , and the sampling fraction is 60%.

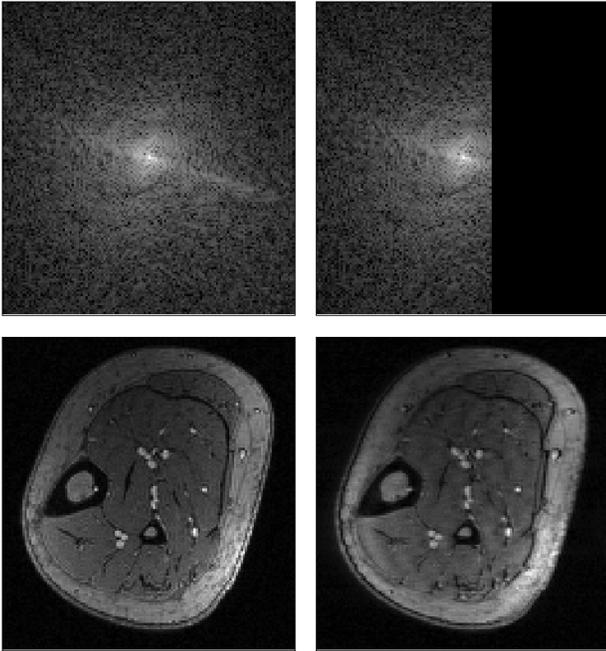


Figure 2. Top row: Fully sampled and PF-sampled k-space from the same source data. Bottom row: Inverse Fourier transform of the top row. The bottom right image is blurred laterally as a result of the PF undersampling.

problem by solving a convex optimization problem. The method is evaluated on complex-valued images generated from both natural imaging and genuine MRI data.

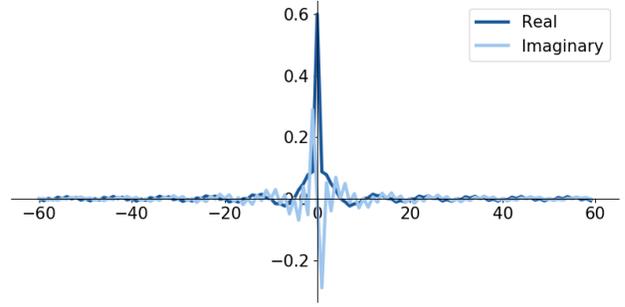


Figure 3. Cross-section of the point spread function (PSF) for 60% PF sampling. The PSF is constant along the fully sampled dimensions. The horizontal axis is in units of pixels.

3.1. Data Pre-Processing

A variety of images will be examined to elucidate the strengths and weaknesses of various priors. It is especially helpful to have a few examples with higher signal-to-noise ratio (SNR) and simpler structure than can be found in MRI. To this end, I make use of a few natural RGB images. The first two color channels are reassigned to the real and imaginary parts of the image, and the blue channel is discarded in order to simulate a complex-valued image.

All source images used in this work have k-space which is fully sampled (i.e. the image is not blurred). All complex-valued source images are normalized to have magnitude no greater than 1. These images serve as the ground truth reference for all experiments. PF-sampled image data is generated retrospectively by masking in the k-space domain (point-wise multiplication with the PF mask in Figure 1).

To simulate real MRI noise statistics, additive white gaussian noise (AWGN) with standard deviation 0.1 is added to the image in k-space immediately before the PF sampling mask is applied. Thus image noise is also gaussian, but is colored as a result of the PF sampling. From another perspective, the PSF blurs the image noise so that the noise at nearby pixels is correlated.

3.2. Optimization Objective

Partial Fourier image formation is a linear transformation, which can be represented as:

$$Cx = b \quad (1)$$

Here C is a circulant matrix which blurs the fully sampled image x with the PF PSF (Figure 3) along the partially sampled dimension. Variable b is the blurred image we get after Fourier transforming PF k-space.

Recovering x from b by simply inverting C is an ill-posed problem, meaning there are many possible solutions. This

is most apparent in k-space, considering that the OTF has many zeroes. I regularize the problem by constructing an objective function which includes not only a data consistency term, but also a weighted prior.

$$\operatorname{argmin}_x \|Cx - b\|_2^2 + \lambda g(z) \quad \text{s.t.} \quad z = Ax \quad (2)$$

3.3. Image Priors

I investigate three image priors in this work. All are convex. The first is total variation (TV), on the basis that it is perhaps the simplest, most successful prior widely used today. This prior promotes sparse gradients, which works well for images that are approximately piece-wise constant. Second, I investigate a Hessian-based prior as a second-order evolution of the TV prior. This prior has been shown to outperform TV in similar non-blind deconvolution problems with biological images [6], [4]. The associated $g(z)$ and A matrices for these priors are described below. Note that total variation has two variants. Anisotropic is simpler to derive and implement, but slightly biased with respect to gradient orientation.

Anisotropic Total Variation (TV_a):

$$A = \begin{bmatrix} \nabla_x \\ \nabla_y \end{bmatrix} \quad (3)$$

$$g(z) = \|z\|_1 \quad (4)$$

Isotropic Total Variation (TV_i):

$$A = \begin{bmatrix} \nabla_x \\ \nabla_y \end{bmatrix} \quad (5)$$

$$g(z) = \|z\|_{2,1} \quad (6)$$

Frobenius Norm of Hessian (FH):

$$A = \begin{bmatrix} \nabla_{xx} \\ \nabla_{yy} \\ \nabla_{xy} \sqrt{2} \end{bmatrix} \quad (7)$$

$$g(z) = \|z\|_{2,1} \quad (8)$$

The 2,1 norm requires some clarification. This operator computes the vector L2-norm at each pixel index (along the difference operator dimension) and sums them all together. This is also called a sum-of-norms or a group lasso penalty [9].

3.4. Alternating Direction Method of Multipliers

I use the Alternating Direction Method of Multipliers (ADMM) to minimize the convex objective (equation 2).

For more details on this algorithm, the interested reader is referred to [1]. The update equations are as follows:

$$x^{k+1} = (C^H C + \rho A^T A)^{-1} (C^H b + \rho A^T (z^k - u^k)) \quad (9)$$

$$z^{k+1} = S_{\lambda/\rho}(Ax^{k+1} + u^k) \quad (10)$$

$$u^{k+1} = u^k + Ax^{k+1} - z^{k+1} \quad (11)$$

The particular form of the soft thresholding function S depends on the prior used. For TV_a, the element-wise soft thresholding operator is used. The block-wise soft thresholding operator is used for TV_i and FH. The reader is referred to [9] for complete details on these proximal operators.

I iterate through these update equations until the residual converges to within 0.1% of its current value. This usually takes less than 50 iterations. The ADMM convergence parameter ρ is set to 1 for all experiments. The regularization weight λ was optimized using a simple grid search.

3.5. Evaluation

The proposed method is evaluated with each prior quantitatively and qualitatively. The image similarity metric PSNR is used to compute the similarity between the magnitude of the complex image output and the magnitude of the ground truth image. Qualitative inspection is simply done by eye.

By construction, the source data serves as a perfect ground truth for comparison. The most appropriate alternative method to compare to is actually quite simply the PF-blurred input image. Other established methods for PF reconstruction [8], [3] will actually degrade image quality in this work because the images do not have slowly-varying phase (as mentioned in the Related Work section).

4. Results

Deconvolution results are illustrated for the four test images in Figures 4-7. In all cases, ADMM converges approximately 30 iterations. Quantitative results are provided in the form of image similarity metric PSNR in Table 1.

For the Mondrian art image (Figure 4), TV performs the best both qualitatively and quantitatively, with little difference (less than 1dB) between the anisotropic and isotropic variants. TV is expected to perform very well here because the image is piece-wise constant. Moreover, the two TV variants should perform very similarly because nearly all gradient information should be parallel to one of x or y.

For the butterfly image (Figure 5), now TV_i and FH perform the best, with less than 1dB difference between the two. The perceptual quality difference between these two outputs is negligible.

| Image | Input | TVa | TVi | FH |
|------------|-------|------|------|------|
| Mondrian | 22.0 | 33.5 | 32.7 | 26.1 |
| Butterfly | 22.8 | 23.8 | 26.2 | 25.3 |
| MRI Leg | 22.4 | 24.1 | 25.9 | 25.2 |
| MRI Pelvis | 24.1 | 27.0 | 28.5 | 28.1 |

Table 1. Quantitative image similarity of ADMM output. Listed values are PSNR (in dB) of the magnitude image as compared with the fully sampled ground truth magnitude image.

The two MRI images (Figures 6 and 7) present the most interesting results. They follow the same regularization ranking as the butterfly image (TVi, then FH, then TVa), but the PSNR difference between TVi and FH is even smaller than for the butterfly. It is the author’s opinion that the FH-regularized images appear more correct, despite the slight lag in PSNR.

The PSNR values in Table 1 were consistent to within ± 0.5 dB between trials. All experiments were run again at $\rho = 0.1$ and $\rho = 10$ to check sensitivity to this parameter. All PSNR values in Table 1 were maintained to within ± 0.5 dB.

5. Discussion

5.1. Choice of Priors

The TV and FH priors were selected on the basis of precedent, but also because they are convex and simple enough to easily derive closed form update equations. Intuitively, we should be able to incorporate more prior knowledge than simply the notion of piece-wise constancy or smoothness. The challenge lies in the practical implementation and optimization of an objective which these new priors.

One source of additional prior information can be motivated by traditional PF reconstruction methods. These techniques greatly leverage Hermitian symmetry in k-space, but only work well on almost real-valued images. Future work could investigate some way to relax this notion of Hermitian symmetry for broader applicability. For example, an image could be regarded as the sum of a real and complex image, the former having Hermitian symmetry, and the latter some desirable characteristic such as sparsity.

Future work could also exploit the fact that we have a univariate PSF operating on multivariate image data. In other words, blurring is technically limited to a single dimension. If an image presents some degree of self-similarity, particularly in orthogonal directions, then it may be possible to learn some features or statistics along the non-blurred image dimensions. Though some caution is warranted here: non-blurred dimensions are still corrupted in some sense by the blurring.

5.2. k-space Filling

The deconvolution methods presented here operate exclusively in the image domain, but there are nevertheless a few interesting observations to be made in k-space. It is clear from all four image examples that the TV prior has a tendency to fill much more of the missing k-space area than the FH prior. And between the two TV variants, the anisotropic version appears to fill even more of k-space. Intuitively, it makes sense that the FH prior would fill in less high-frequencies than TV since if piece-wise smoothness is indeed a more fitting image attribute than piece-wise constancy.

5.3. Problem Construction

Data consistency is performed in the image domain out of computational convenience. It allows for both A and C matrices to be circulant so that they can be efficiently computed by Fourier multiplication. The data consistency term is MAP-optimal because the L2 norm is invariant under the Fourier Transform (by Parseval’s theorem) and we have AWGN in k-space.

5.4. Data Dimensionality

I chose to work with 2D image for speed of computation and simplicity. Future work could build upon this by investigating the potential advantage of running this algorithm on 3D data. After all, adding a 3rd dimension means adding a second non-blurred dimension. It seems plausible that this additional spatial context could improve the performance of priors which promote some order of spatial continuity, such as TV and FH.

5.5. Non-Local Means

The orthogonal self-similarity prior described above also motivates an entirely different approach to solving this deconvolution problem. With homage to the popular non-local means algorithm for denoising [2], similar pixel patches could serve as examples from which to average and replace some target pixel. Again, this relies on the claim that non-blurred dimensions retain some semblance to the

Unlike the traditional non-local means, the patches would need to be 1D strips of pixels oriented orthogonal to the target pixel, and a comparison would be made between the target pixel strip and the blurred example strips. I implemented a basic version of this algorithm and found it to perform very well on a few toy examples, but such an algorithm is highly unscalable due to repetitive exhaustive searching. The brute-force implementation is $O(N^2)$ to iterate once on an image with N pixels.

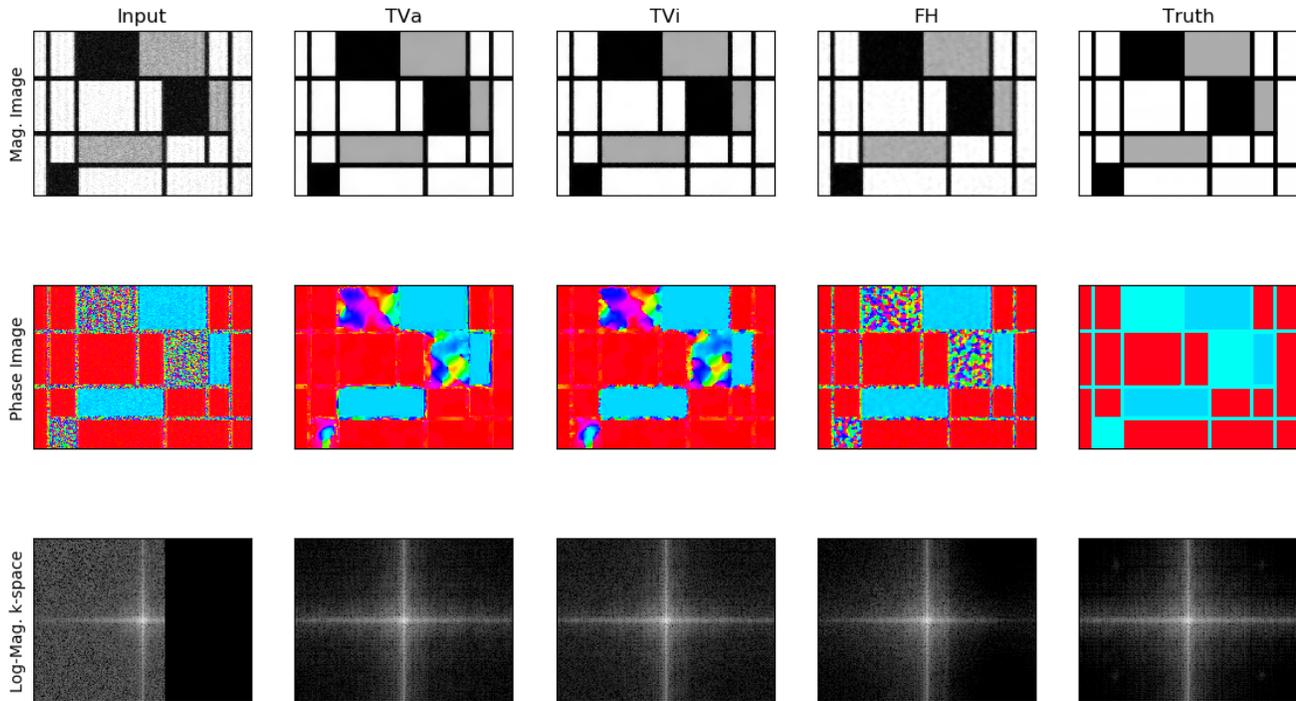


Figure 4. Deconvolution results for the simplest test case, a Mondrian artwork. $\lambda = 0.09, 0.09, 0.03$ is used as the regularization weight for TVa, TVi, and FH respectively.

References

- [1] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *MAL*, 3(1):1–122, July 2011.
- [2] A. Buades, B. Coll, and J. Morel. A non-local algorithm for image denoising. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pages 60–65 vol. 2, June 2005.
- [3] E. Haacke, E. Lindskog, and W. Lin. A fast, iterative, partial-fourier technique capable of local phase recovery. *Journal of Magnetic Resonance (1969)*, 92(1):126–145, Mar. 1991.
- [4] H. Ikoma, M. Broxton, T. Kudo, and G. Wetzstein. A convex 3d deconvolution algorithm for low photon count fluorescence imaging. *Scientific Reports*, 8(1):11489, July 2018.
- [5] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser. Deep Convolutional Neural Network for Inverse Problems in Imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522, Sept. 2017.
- [6] S. Lefkimiatis, A. Bourquard, and M. Unser. Hessian-Based Norm Regularization for Image Restoration With Biomedical Applications. *IEEE Transactions on Image Processing*, 21(3):983–995, Mar. 2012.
- [7] G. McGibney, M. R. Smith, S. T. Nichols, and A. Crawley. Quantitative evaluation of several partial fourier reconstruction algorithms used in mri. *Magnetic Resonance in Medicine*, 30(1):51–59, 1993.
- [8] D. C. Noll, D. G. Nishimura, and A. Macovski. Homodyne detection in magnetic resonance imaging. *IEEE Transactions on Medical Imaging*, 10(2):154–163, June 1991.
- [9] N. Parikh and S. Boyd. Proximal Algorithms. *OPT*, 1(3):127–239, Jan. 2014.
- [10] J. Schlemper, J. Caballero, J. V. Hajnal, A. Price, and D. Rueckert. A Deep Cascade of Convolutional Neural Networks for MR Image Reconstruction. In *Information Processing in Medical Imaging*, Lecture Notes in Computer Science, pages 647–658. Springer, Cham, June 2017.
- [11] S. Wang, Z. Su, L. Ying, X. Peng, S. Zhu, F. Liang, D. Feng, and D. Liang. Accelerating magnetic resonance imaging via deep learning. In *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*, pages 514–517, Apr. 2016.

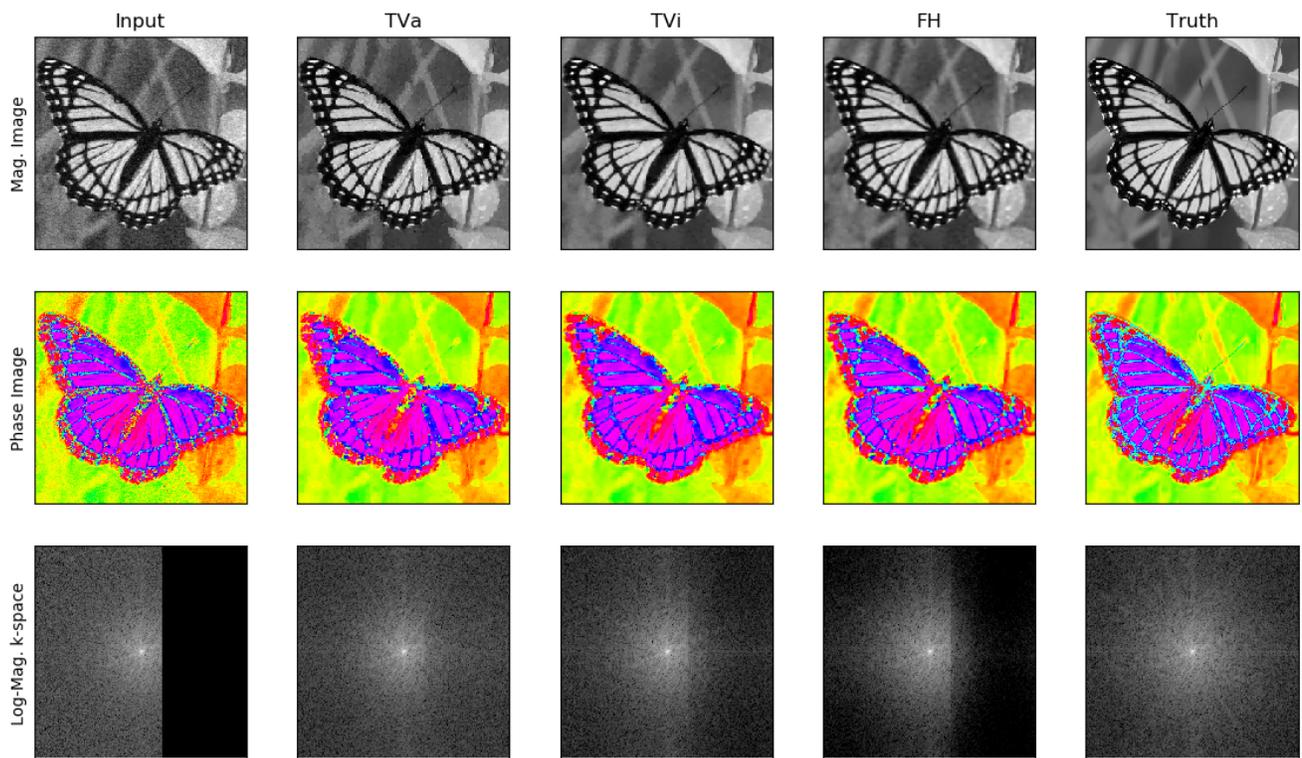


Figure 5. Deconvolution results for a natural image of a butterfly. $\lambda = 0.04, 0.05, 0.03$ is used as the regularization weight for TVa, TVi, and FH respectively.

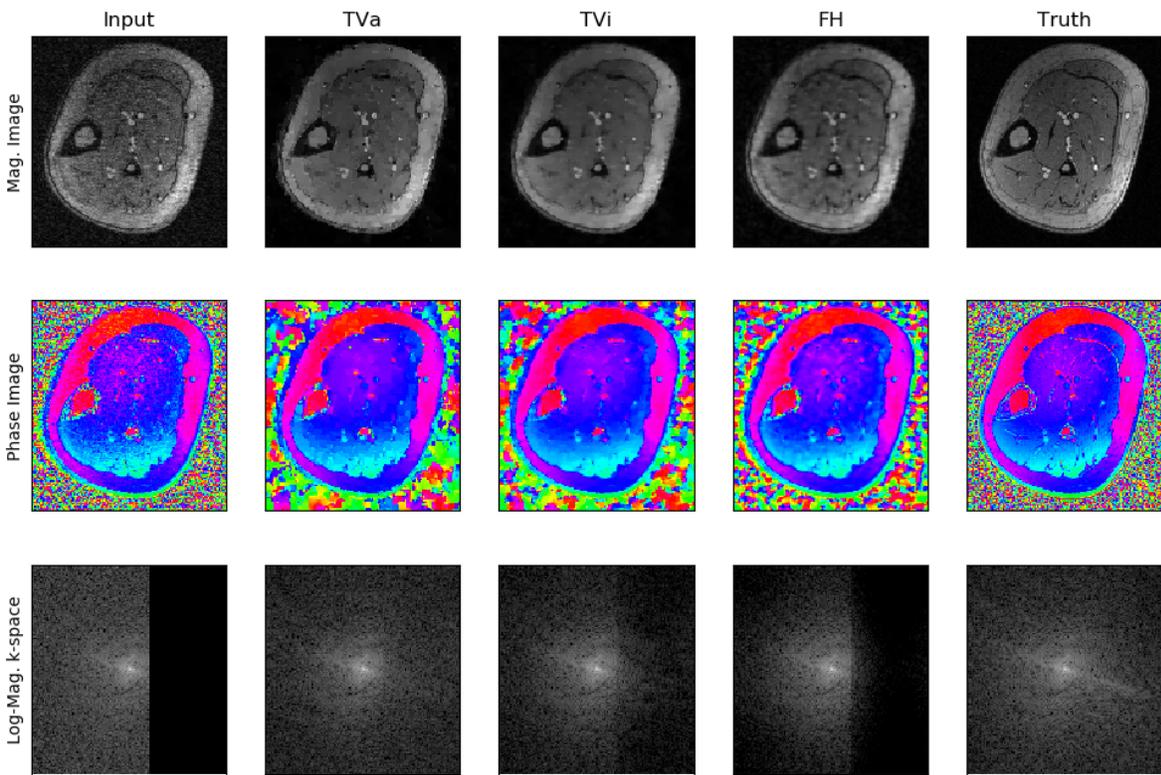


Figure 6. Deconvolution results for a transverse plane MRI of a human leg. $\lambda = 0.04, 0.05, 0.02$ is used as the regularization weight for TVa, TVi, and FH respectively.

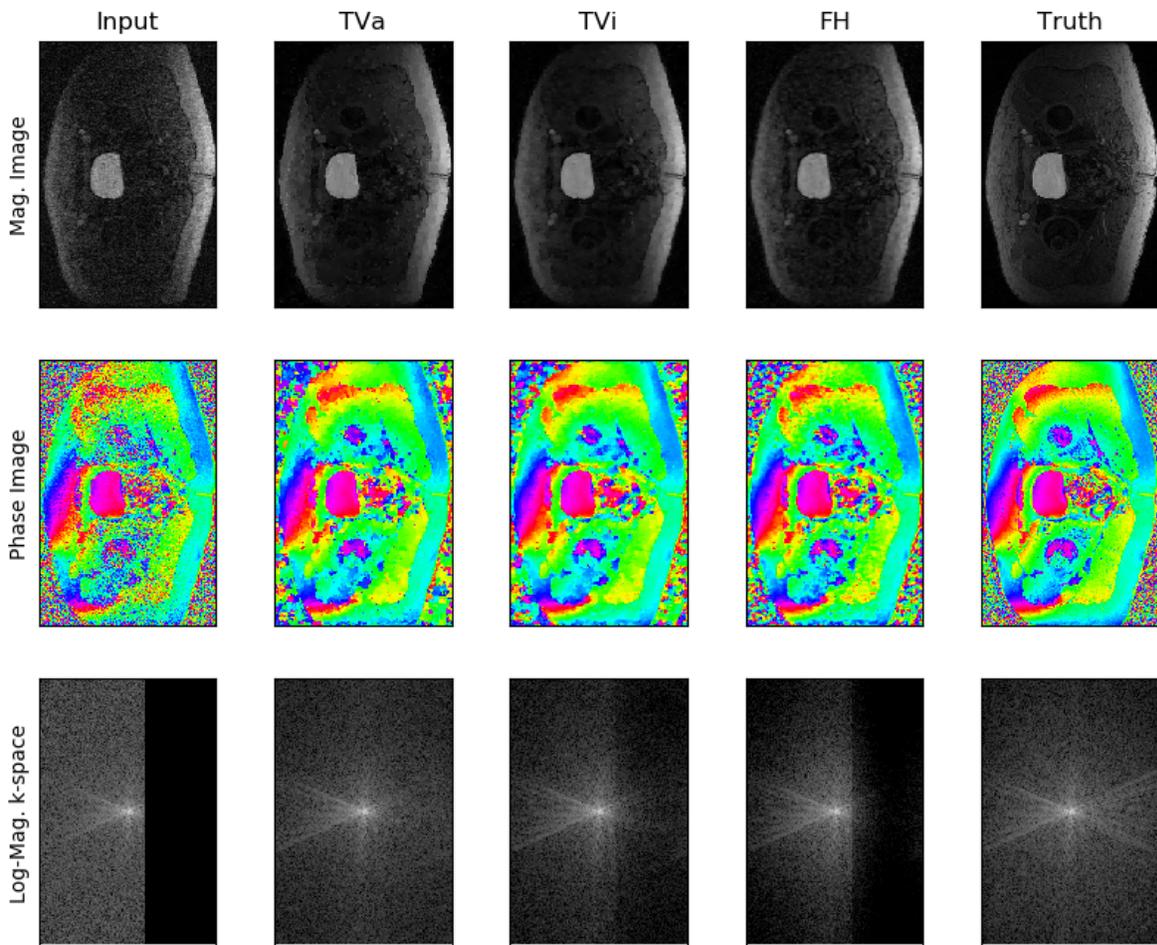


Figure 7. Deconvolution results for a transverse plane MRI of a human pelvis. $\lambda = 0.04, 0.05, 0.02$ is used as the regularization weight for TVa, TVi, and FH respectively.