Image Fusion for MR Bias Correction

Ayres Fan

Stochastic Systems Group

Joint work with W. Wells, J. Fisher, M. Cetin, S. Haker, A. Willsky, B. Mulkern

Magnetic Resonance

The magnetic resonance (MR) effect can be observed in a variety of settings. For our purposes, we are only interested in the hydrogen atom. Hydrogen is useful for human imaging because the body is 70% water.

A large homogeneous magnetic field (B0) is applied in the +z direction. This causes all protons to align their spins either in the +z or -z direction. The proton spin is not directly aligned with z but instead precesses around that axis. The key is the Larmour equation which states that the resonant frequency of the precession is proportional to the magnitude of B0. Thus if we apply another field on top of B0, one that varies in space, we can vary the resonant frequency with space.

We can then apply a RF pulse to the object in question in order to observe the resonant response. The response of the protons as a function of space will then be encoded in the frequency response of the system.

MR Imaging

MR imaging is widely used in medical applications for its flexibility and excellent soft tissue contrast. The image generation is controlled by three intrinsic tissue properties and two user definable parameters (assuming the pulse sequence has already been decided):

$$I = \rho \exp(-\mathsf{TE}/\mathsf{T2})(1 - \exp(-\mathsf{TR}/\mathsf{T1}))$$

TE: time echo (time we measure signal)

TR: time repeat (time between pulse sequences)

T1: spin-lattice relaxation (recovery of z-magnetization)

T2: spin-spin relaxation (loss of xy-magnetization)

ρ: proton density

We can target T1 and T2 through appropriate selection of TE and TR.

Image Reconstruction

The MR physics people often term the frequency domain k-space. The noise in the imaging process is thermal and can be adequately modeled as Gaussian. We can convert from k-space to real space through an inverse Fourier transform. Because this is a linear transformation, the noise remains Gaussian.

Unfortunately due to noise and such effects, the result of the IFFT will be complex. The most common protocol is to represent the final image as the absolute value of the complex image. This will change the noise from Gaussian to Rician.

With Rician noise, the noise depends on the signal-to-noise ratio (SNR). At low SNR, the noise is approximately Rayleigh. At high SNR, the noise is approximately Gaussian. In general, the Rician pdf is messy and difficult to work with.

Bias Field

Typically the transmitting coil is a body coil. The most important design feature is spatial homogeneity of the induced magnetic field. The receiving coil is often a surface coil which is placed close to the surface of the object to be imaged. The most important feature is high frequency response at the Larmour frequency (high Q ratio).

The signal observed at the receiver is then:

$$I(x) = \rho(x) b(x) \exp(-\mathrm{TE}/\mathrm{T2}(\mathrm{x}))(1 - \exp(-\mathrm{TR}/\mathrm{T1}(\mathrm{x})))$$

b(x) is simply the magnetic field induced by the receiving coil. If the exact coil location is known, b(x) can be calculated using the Biot-Savart Law. For surface coils, b tends to be large near the coil and falls off as 1/r.

We can then write a model for our observed image as:

$$I(x) = b(x) f(x) + n(x)$$

Previous Work

The earliest work dealing with the MR bias problem occurred in the mid-80s with work by people such as Axel et al (1984) on using homomorphic unsharp filtering. This relies on the assumption that the bias field and the MR image are separable in the image domain. Other work at the time involved correcting images with bias fields obtained from phantoms.

More recent work has included work by Dawant et al (1993) where they have users select points within a region that is homogeneous except for the bias field. They then fit parametric functions to the points.

A popular approach was that of Wells et al (1996). They interleaved bias estimation and a statistical classifier to both correct and segment MR images.

The problem with previous approaches is that they either don't work well, require a great deal of user supervision, or they require data that can be segmented using a statistical classifier.

Measurement Model

In 1988, Brey and Narayana proposed an approach that involved capturing images from both the body coil and the surface coil. The image from the body coil will be homogeneous (in terms of the induced magnetic response) while the image from the surface coil will have good SNR in the region of interest (ROI).

Our measurement model is then:

$$\begin{aligned} I_B(x) &= kf(x) + n_B(x) \\ I_S(x) &= b(x)f(x) + n_S(x) \end{aligned}$$

They estimate b as:

$$\widehat{b} = (h_1 * I_S)/(h_2 * I_B)$$

and they estimate f as:

$$\widehat{f} = (h_3 * I_S)/\widehat{b}$$

ML Formulation

The idea of using this side information from the body coil scan seems promising, though the Brey-Narayana method seems a bit ad hoc.

For simplicity, we will model the noise as Gaussian and iid. This is approximately valid in regions of high SNR and will create suboptimal estimates in regions of low SNR.

We convert our 2D or 3D images into vectors and stack them together to form one large observation vector:

$$I = \begin{pmatrix} I_B \\ I_S \end{pmatrix}$$

We also represent b and f as vectors. We can write the log likelihood as follows (ignoring constant terms):

$$l(I; b, f) = -\frac{1}{2\sigma_B^2}||I_B - kf||^2 - \frac{1}{2\sigma_S^2}||I_S - b \circ f||^2$$

Regularization

If we attempt to maximize that function, we find that the answer we obtain is useless:

$$\begin{aligned} \hat{b} &= I_S/I_B \\ \hat{f} &= I_B \end{aligned}$$

So we can add regularizers to f and b to produce more interesting results. We can write a total energy functional that we wish to minimize:

$$E(b,f) = ||I_B - kf||^2 + \lambda||I_S - b \circ f||^2 + \alpha\mathcal{R}(b) + \gamma\mathcal{R}(f)$$

Generally the regularizing functions will be high pass. The bias field is smooth (because it's an induced magnetic field, the solution must satisfy Laplace's equation). The true underlying MR signal tends to be piecewise constant within tissue boundaries.

Solution for f

We set γ=0 and choose a quadratic regularizing function for b. A quadratic regularizer for f does not make sense (due to the edges that are present), and non-quadratic solvers are much slower and generally are more difficult to solve. Eventually we would like to try a total variation regularizer for f.

We minimize E using coordinate descent, alternating minimizations on f and b. Minimization in terms of each parameter is easier than minimizing with respect to both at the same time.

This restriction makes the solution for f trivial. There is no spatial coupling in f, so we can solve for f on a pointwise basis:

$$\hat{f}_n = \frac{kI_B(n) + \lambda b_n I_S(n)}{b_n^2 + k^2}$$

Rewriting the Energy Functional

We choose regularizing functions for b of the form:

$$\mathcal{R}(b) = ||Lb||^2$$

We mainly work with L matrices that implement derivative operators (such as gradient or the Laplacian) though arbitrary high-pass filters are acceptable.

We define a diagonal matrix F that has the entries of f on the main diagonal. We can then rewrite our energy functional as:

$$E(f, b) = \frac{1}{2} b^T Q b - a^T b$$

where Q is a PSD matrix:

$$Q = 2\lambda F^2 + 2\alpha L^T L$$

and a is a vector:

$$a = 2\lambda F I_S$$

Quadratic Optimization

We can find the minimum with respect to b exactly, though inverting a 65536x65536 matrix is quite slow. There are a number of suboptimal iterative schemes that can be used to find approximations to the minimum. These include gradient descent and conjugate gradient. Conjugate gradient is easy to implement because exact line minimizations can be conducted in O(n) time.

We can also apply a preconditioner to Q to reduce the condition number. In the end, we use a tridiagonal preconditioner with the entries coming from the three main diagonals of Q. A quindiagonal preconditioner was tested but did not significantly improve convergence speed.

Parameters

There are a large number of parameters to set. The noise variances can be estimated directly from the image in homogeneous regions (the best are air-filled regions such as the rectum). The scaling parameter k can be controlled during the imaging process.

One interesting parameter is the choice of how long to let the iterative solver on b run. If we knew that our f estimate was correct, we would want to let the b solver run to completion. But f is not necessarily correct so by letting the b solver run for too long, we could in fact be moving away from the global minimum.

We observed this phenomenon when choosing to let the b solver run for more iterations actually led to more overall coordinate descent iterations.

Initialization

A wide variety of initial guesses can be chosen for f and b. In general this choice does not affect whether we get a correct answer, but it can play a significant role in convergence speed.

We found that the best initialization comes from using Brey and Narayana's method. The b estimate from homomorphic unsharp filtering was found to be vastly inferior.

Convergence

There is not a great deal we can say about the convergence other than our empirical observations. We are guaranteed to converge to a local minimum. This can be seen from the fact that each f or b iteration is guaranteed to decrease the energy functional. On the f step, we find the exact minimum. On the b step, each conjugate gradient step is a descent step.

The energy functional is not convex but we do not have problems with local minima. Even when we start with initial f and b estimates as random (but positive), we still converge to the correct answer.

Example Data

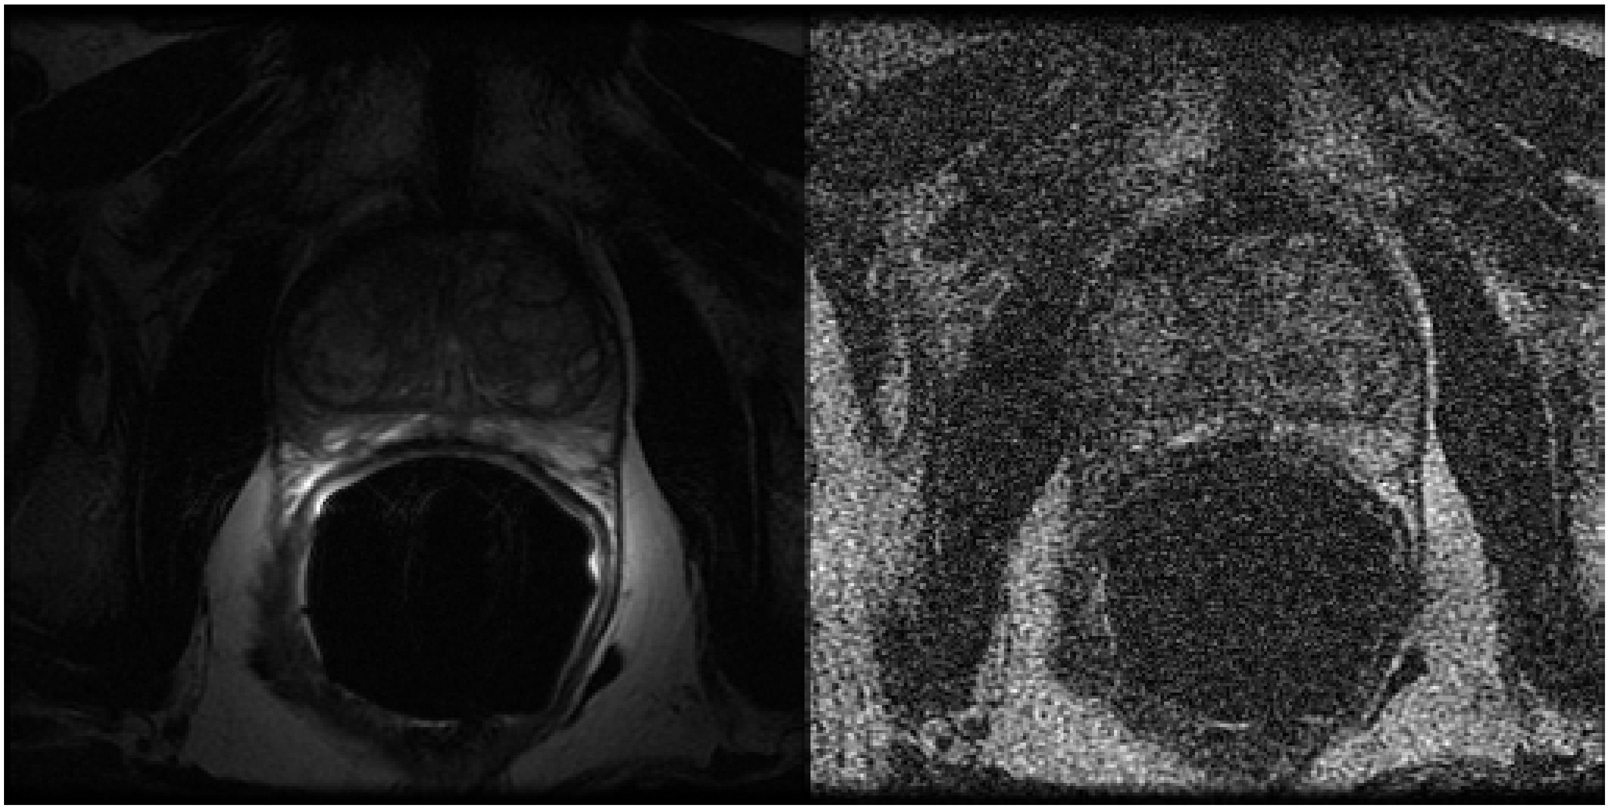

Bias Corrected Images

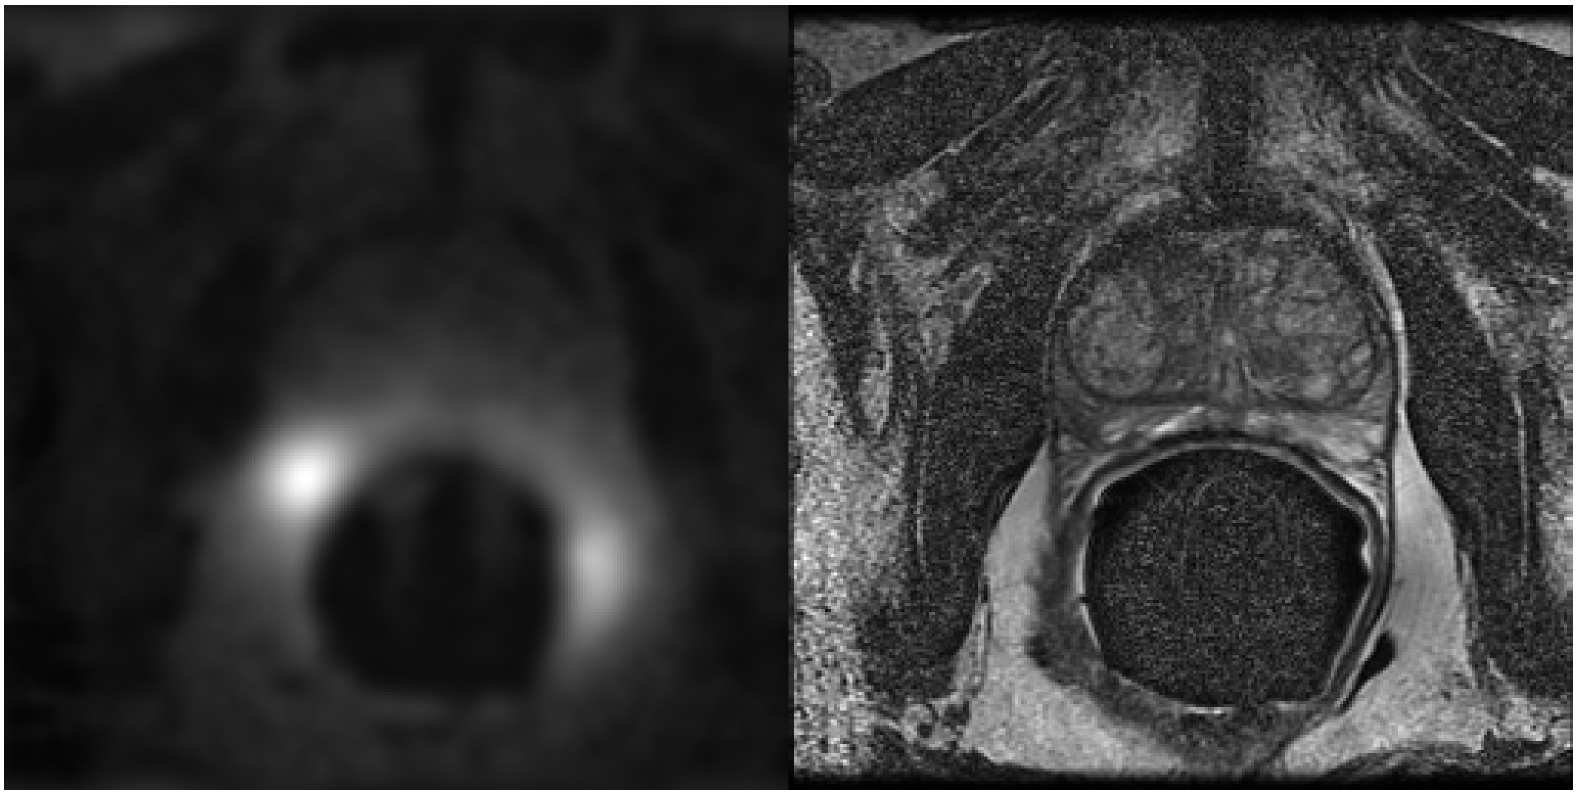

Varying alpha

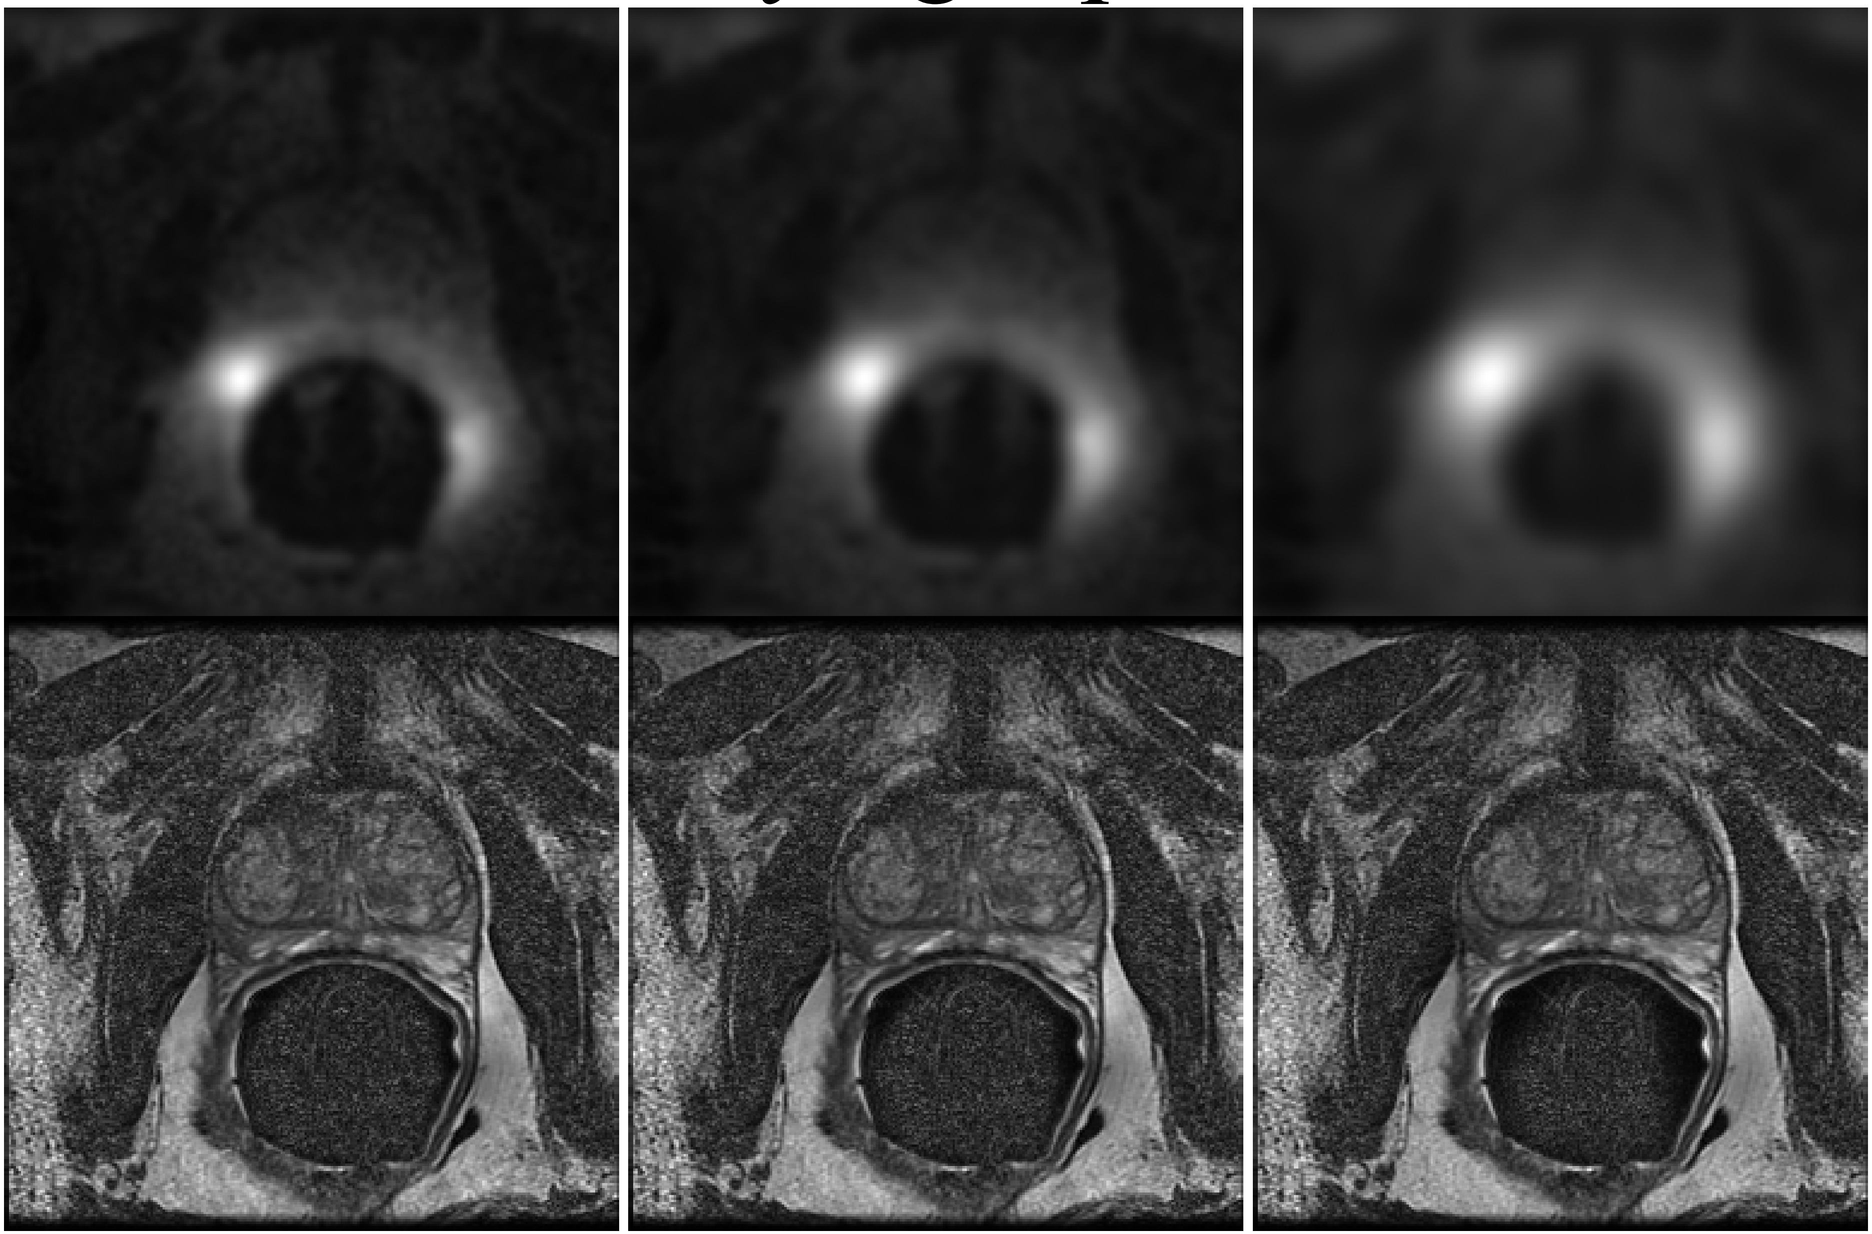

Another data set

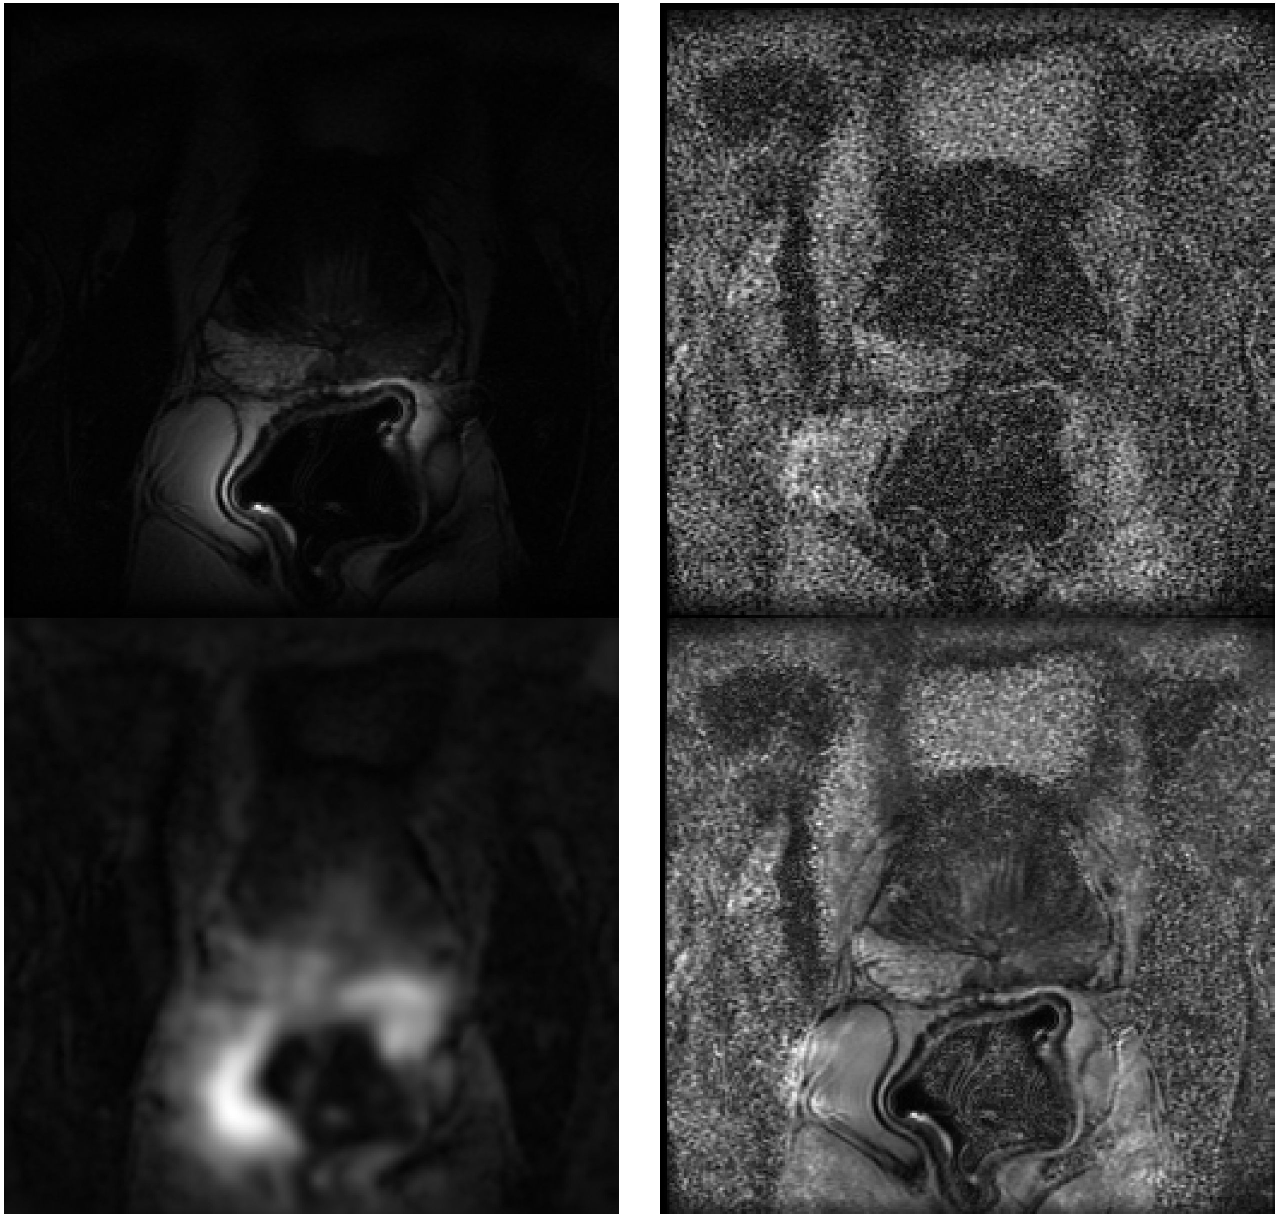

Applying correction to T1

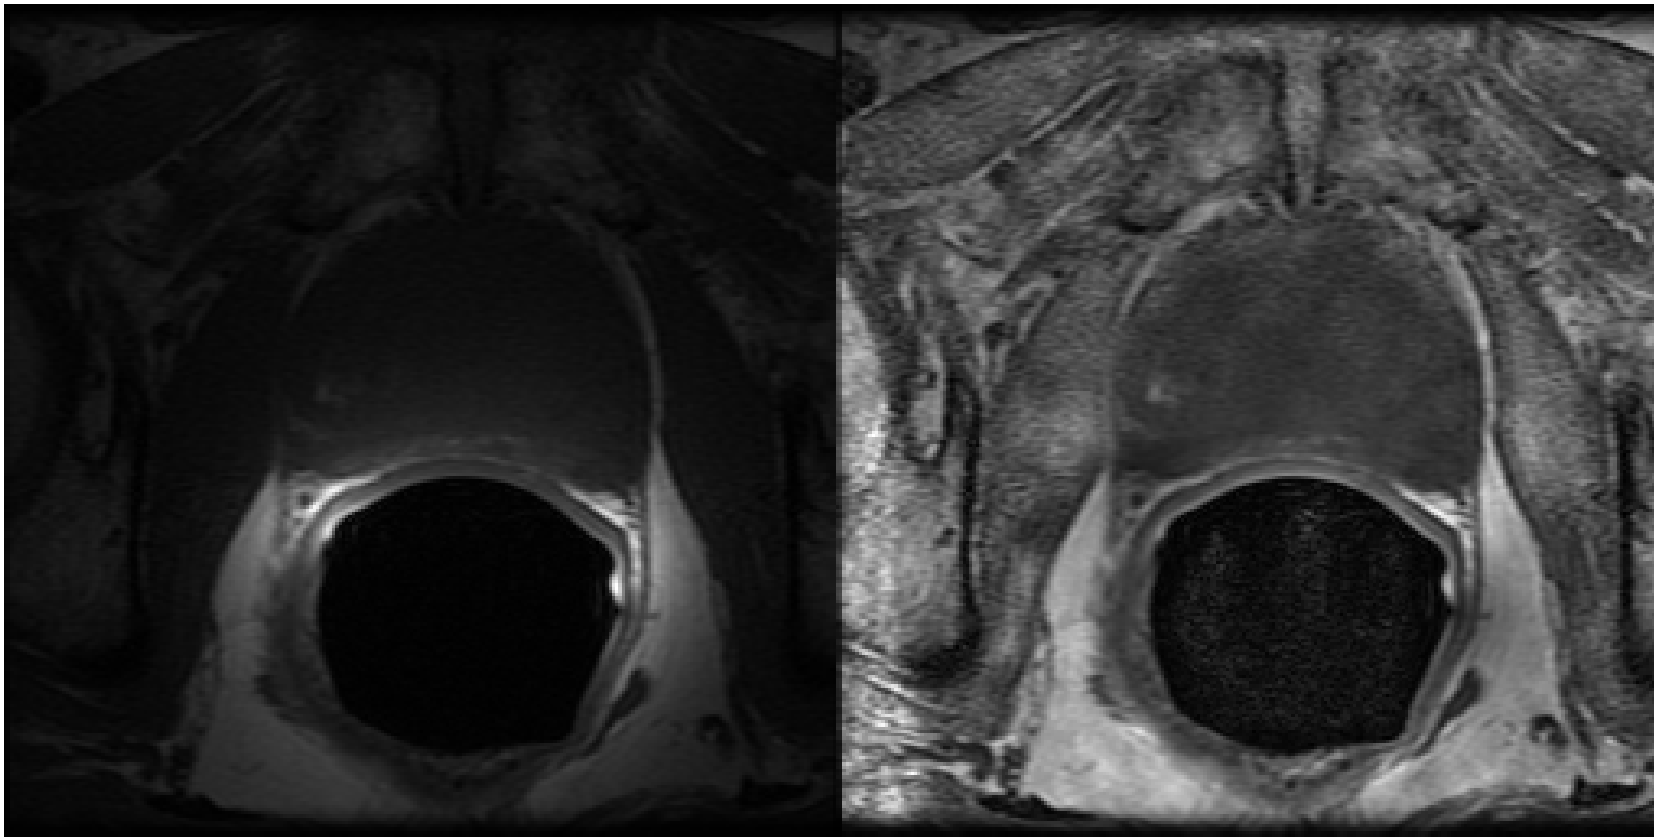

Gradient fields

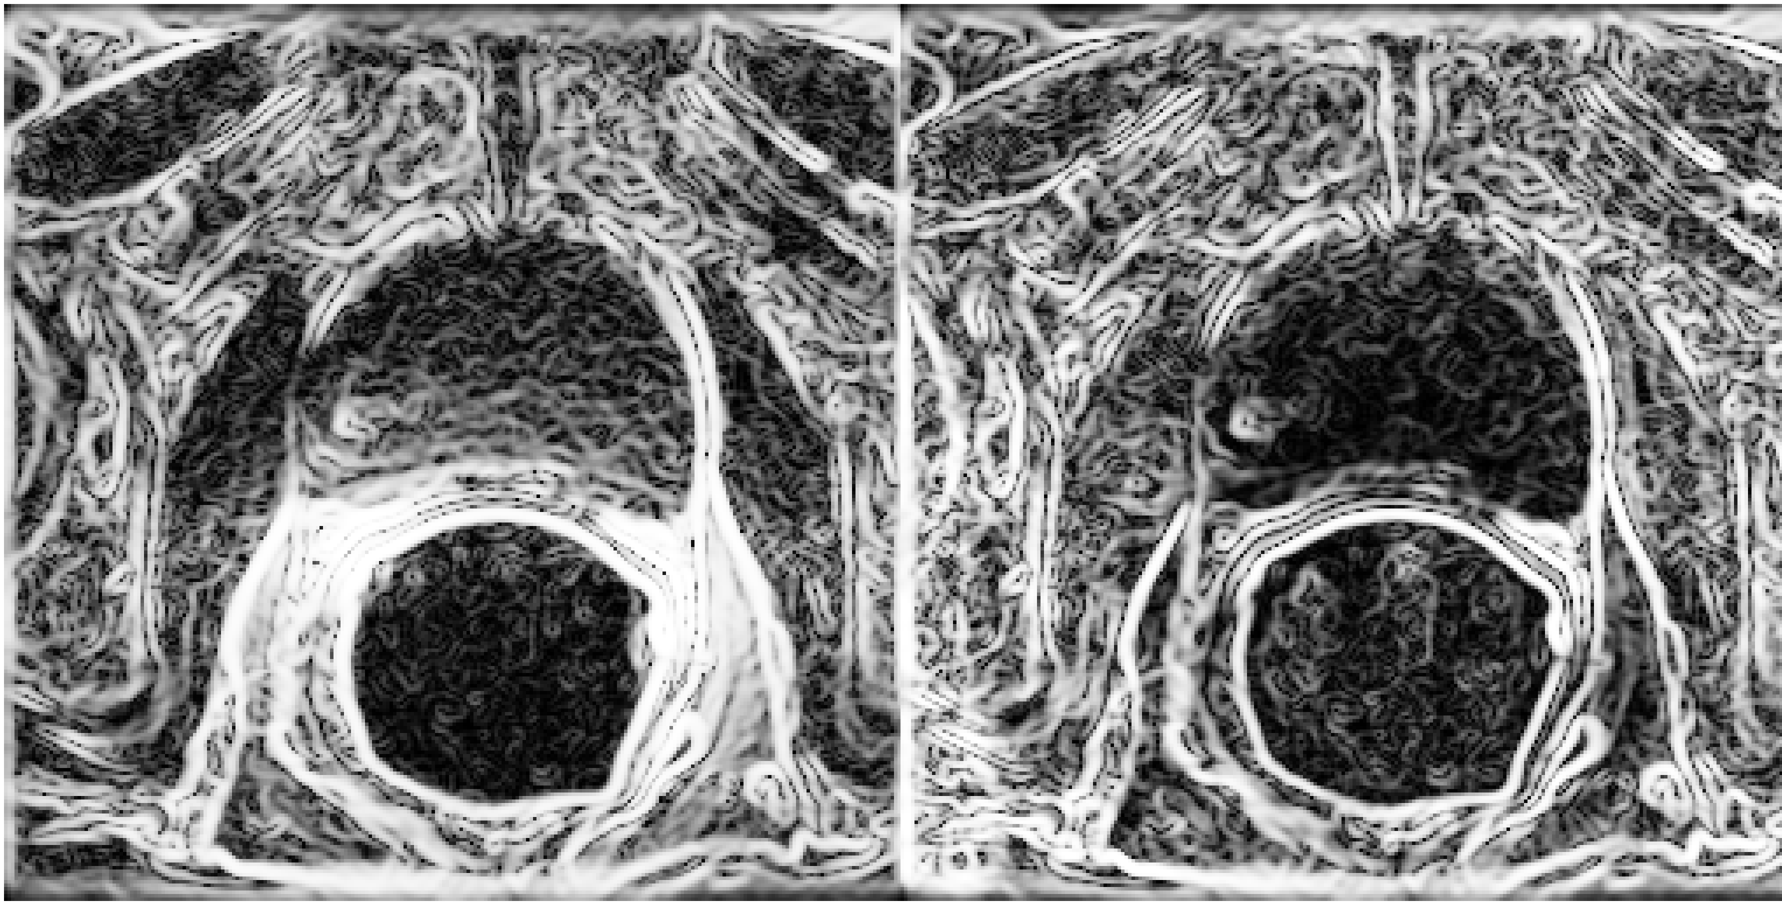

Future Directions

- Couple with registration
- Try on more data
- 3D
- Different regularizers
- Multiresolution