MR Bias Correction and Reflectance and Illumination Separation

Ayres Fan

Stochastic Systems Group

Massachusetts Institute of Technology

September 24, 2003

Outline

Basic Formulation

- We have two multiplicative fields with additive noise:

$$y(x) = f_1(x)f_2(x) + n(x)$$

- We wish to estimate both f_1 and f_2 using statistical knowledge of the noise as well as prior information on the fields

MR Problem Statement

- The bias field is a systematic intensity inhomogeneity that corrupts magnetic resonance (MR) images.

- Correcting for the bias field makes both human analysis (e.g., tumor detection, cartilage damage assessment) and computer analysis easier (e.g., segmentation, registration).

- General assumptions:
  - The bias field is slowly varying in space.
  - The bias field is tissue independent.
  - Tissue intensities are piecewise constant.
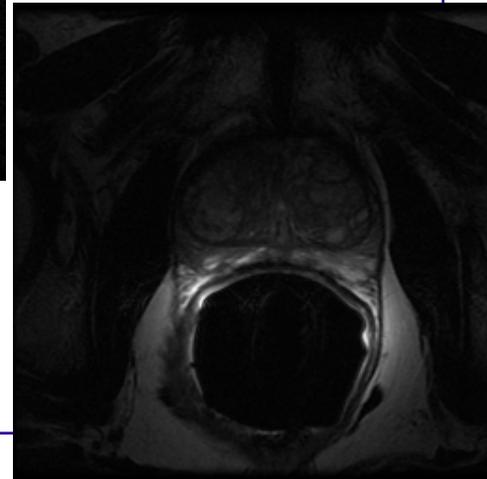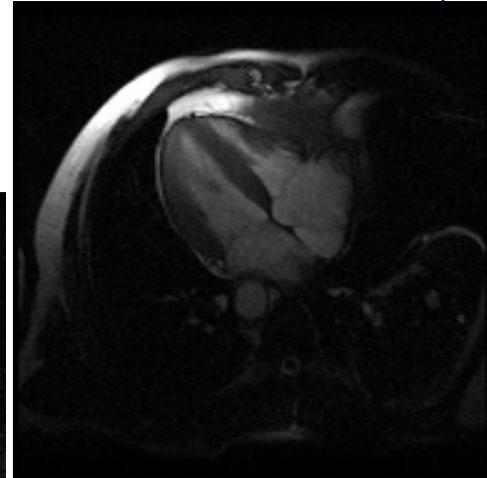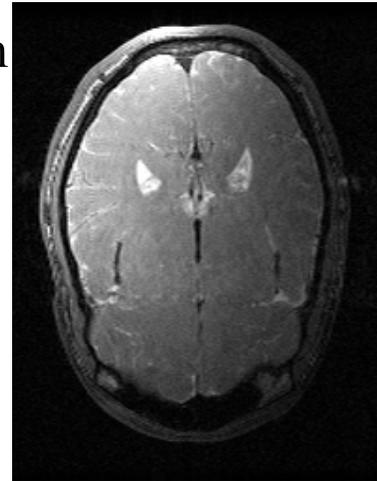
Illumination and Reflectance

- We want to try to separate the image into illumination and reflectance components as best as possible

- Reflectance map will contain albedo and texture

- Multiplicative model breaks down with non-Lambertian reflectance models (e.g., when there are specularities)

- Both maps will have edges

Outline

MR Imaging

- We would ideally want an image that is solely dependent on tissue-dependent parameters and user controlled parameters, e.g.,

$$\varphi(x) = \rho(x)\exp(-T_E/T_2(x))(1-\exp(-T_R/T_1(x)))$$

- We can target ρ, T1, and T2 measurements through appropriate selection of TE and TR.

- Received MR data is in k-space (frequency domain) with Gaussian noise. Use IFFT and take absolute value to obtain reconstructed image. This results in Rician noise.

- Maximizing SNR is a major goal. We can increase SNR by spatial averaging, time averaging, filtering, increasing signal reception, and increasing the strength of the applied magnetic field.

Bias Field

- The body coil is a large coil usually wrapped around the main cavity
  - Typically transmit RF pulse sequence with body coil due to good spatial homogeneity. Transmitted power density limited by FDA.
- Surface coils are coils placed near the object of interest.
  - In order to increase signal level, increase the induced magnetic field strength through coil design.
  - Surface coils have good signal strength near the coil, and the strength rapidly diminishes with distance. Thus we can get good SNR in the ROI.
- The signal observed at the receiver is then:

$$I(x) = \beta(x)\varphi(x) + n(x)$$

  β(x) is the magnetic field induced by the receiving coil.
  φ(x) is the ideal MR signal

- The severity of the bias field is determined by the spatial homogeneity of the surface coil reception profile.

Measurement Model

- Brey and Narayana (1988) proposed capturing images from both the body coil and the surface coil. The measurement model is then:

$$I_B(x) = \varphi(x) + n_B(x)$$
$$I_S(x) = \beta(x)\varphi(x) + n_S(x)$$

- I_B is homogeneous but noisy. I_S has high SNR in the region of interest, but a potentially severe bias artifact.

- Note that gain in SNR from using a surface coil does not come from reduction of noise, but from increased signal gain from the bias field.

Example Data
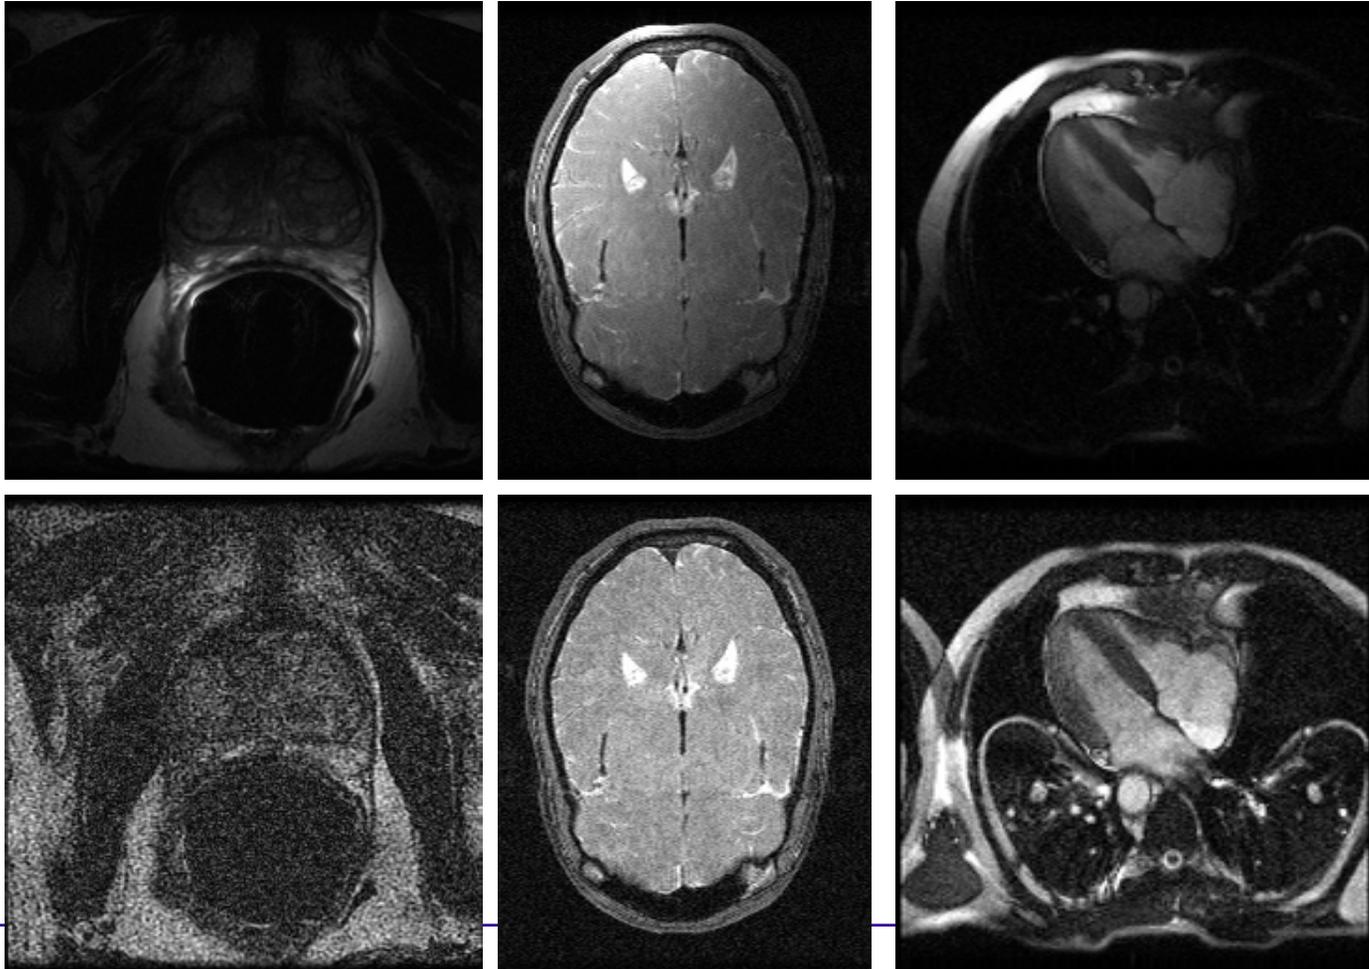
Surface coil images (top) and body coil images (bottom)

Previous BC/SC Work

- Brey-Narayana filter the two observation images to denoise them and divide the results to estimate the bias field:

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

- Lai and Fang (1998) take I_S/I_B and select a sparse set of reliable control points and fit splines to the bias field.

- Pruessmann et al (2001) who fit local polynomials to the bias field.

ML Formulation

- We model the noise as Gaussian and IID.
  - This is a good approximation in medium-to-high SNR regions except the Rician noise adds a bias of 2-5%.
  - Generally low SNR regions correspond to air regions which we do not care about
- We stack the 2D or 3D images into vectors.
  - We now want to estimate the vector quantities **f** and **b**.
  - The discrete measurement model is:

$$y_B = f + n_B$$
$$y_S = b \circ f + n_S$$

- We can then write the log likelihood as (ignoring constant terms):

$$\ell(y_B, y_S; b, f) = -\frac{1}{2\sigma_B^2}||y_B - f||^2 - \frac{1}{2\sigma_S^2}||y_S - b \circ f||^2$$

Regularization

- Finding the ML estimate results in trivial estimates:

$$\hat{b} = y_S/y_B$$
$$\hat{f} = y_B$$

- We construct an augmented energy function that encourages smoothness in **b** and piecewise smoothness in **f**:

$$E(b, f) = \lambda_B ||y_B - f||^2 + \lambda_S ||y_S - b \circ f||^2 + \alpha ||Db||^2 + \gamma ||Lf||_p^p$$

  - We generally choose p <= 1 to help preserve edges
  - **D** and **L** are matrices chosen to implement differential operators
  - λ_B, λ_S, α, and γ are all positive constants
  - The λ's can be seen to be related to the inverse noise variances of the observed images.

Outline

Optimizing Energy Function

- Overall problem is non-convex.
- Use coordinate descent to alternately optimize **b** and **f**
  - Minimizing the energy simultaneously with respect to **b** and **f** is difficult. But given b, f is relatively easy to obtain, and vice versa.
  - A stationary point found using coordinate descent is also a stationary point of the overall energy functional.
- b-step
  - Minimize $\lambda_S \|y_{\mathsf{S}} - Fb\|^2 + \alpha \|Db\|^2$

    F is a diagonal matrix with **f** along the diagonal
- f-step
  - Minimize $\lambda_B \|y_{\mathsf{B}} - f\|^2 + \lambda_S \|y_{\mathsf{S}} - Bf\|^2 + \gamma \|Lf\|_p^p$

    B is a diagonal matrix with **b** along the diagonal

Solving for b

- With **f** fixed, the energy is quadratic in terms of **b**.
  - The quadratic matrix is positive definite which means that the energy function is strictly convex.
  - Hence there is only one local minimum, which is also the global minimum, for a given **f**.

- Setting the gradient to zero leads to a linear equation for the solution:

$$\left(\lambda_S F^2 + \alpha D^T D\right) b = F y_S$$

- We can solve by direct matrix inversion, but **b** is often very large (e.g., 65,536 elements for a 256x256 image).

- We use preconditioned conjugate gradient to find a sub-optimal iterative solution.

Solution for f (γ=0)

- With no regularization on f, we can minimize the function pointwise:

$$\hat{f}[n] = \frac{\lambda_B y_B[n] + \lambda_S b[n] y_S[n]}{\lambda_B + \lambda_S b^2[n]}$$

- This equation can be interpreted as a noise-weighted convex combination between y_B and y_S/b.
  - When **b**[n] is large, we mainly use the surface coil for the reconstruction
  - When **b**[n] is small, we mainly use the body coil
- Increases SNR by 0-3 dB over best image at each point.
  - In regions far from the surface coil, there will be a significant advantage to incorporating the body coil measurements.
- This is essentially the same fusion equation found by Roemer et al (1990). This is the canonical method of combining multiple surface coil images when the coil reception profiles are known.

Half-Quadratic Optimization

- To solve l_p regularization problems, we use half-quadratic optimization (Geman-Reynolds 1992).
  - Fixed-point iterative method where we form a succession of quadratic approximations at $\mathbf{f}^{(0)}$, $\mathbf{f}^{(1)}$, …:

  $$\|Lf\|_p^p \approx f^T L^T W^{(i)} Lf$$

  - Choose $\mathbf{W}^{(i)}$ so that the relationship holds for $\mathbf{f}^{(i-1)}$:

  $$W^{(i)}[n,n] = \frac{p}{2}\left(((D\hat{f}^{(i-1)})[n])^2 + \xi\right)^{p/2-1}$$

- This makes each f-step quadratic, so we just need to solve this linear equation at each half-quadratic iteration:

  $$\left(\lambda_B I + \lambda_S B^2 + \gamma D^T W^{(i)} D\right)\hat{f}^{(i)} = \lambda_B y_\mathsf{B} + \lambda_S B y_\mathsf{S}$$

  - This is again a positive definite system, so we can use PCG to solve.
  - Results in three sets of nested iterations.

Multiple Coils & Pulse Sequences

- Many protocols use multiple surface coils to simultaneously receive the MR signal. This allows for better spatial coverage.

  – Generally combine into a single composite image using sum-of-squares.

- We can generalize our energy functional to find the bias field estimate for each coil and one composite true image estimate.

- When acquiring multiple pulse sequences, we only need one body coil image. The bias field is largely unchanged (though some minor effects may crop up due to, e.g., magnetic susceptibility of the tissue).

- We can use coordinate descent again, but on each b- or f-step we need to compute an estimate for each **b** and each **f**.

3D Processing

- Straightforward application of the exact same energy functional.
  - Slower than independent processing of slices.
  - Enables better coupling across slices than independent processing allows.
- Common clinical practice to capture volumes with different slice orientations (e.g., axial, sagittal).
  - Usually have higher in-plane resolution than inter-plane resolution.
  - Only need one body coil image to correct different orientation volumes because the bias field is slowly changing in space. Simply find the bias field for the volume we have a body coil image for, and interpolate onto other sampling grids.

Optimizing Performance

- Multigrid
  - Multiresolution technique that can help avoid local minima, and increase convergence speed.
  - Simple coarse-to-fine implementation.
  - This allows us to find the low frequency components on coarser grids and propagate the results to the original scale.

- This results in convergence of about 30 sec per slice

- Many other parameters to choose (e.g., tolerances for each f-step and b-step). Need to optimize individually for each particular application.

Outline

1. Introduction to bias correction
2. MR measurement model and problem formulation
3. Iterative Solver
4. *MR Results*
5. Reflectance and Illumination Separation
6. Conclusion

MNI Example

(left to right) True image (top), B-N estimate (bottom), body coil image (top), our f estimate (bottom), surface coil images, and bias field estimates.
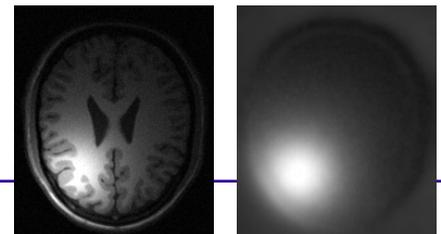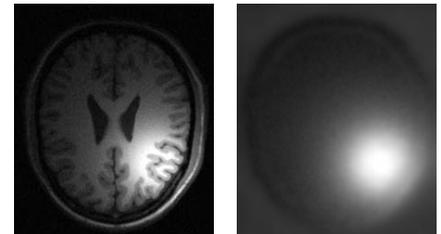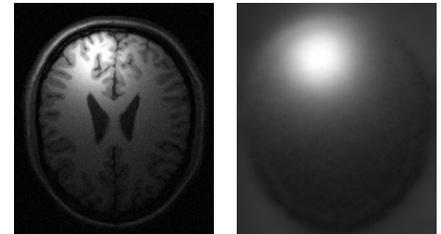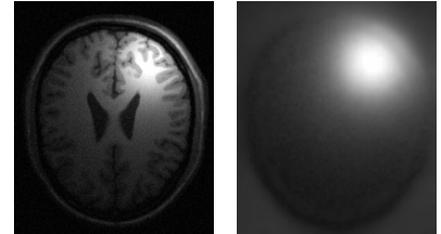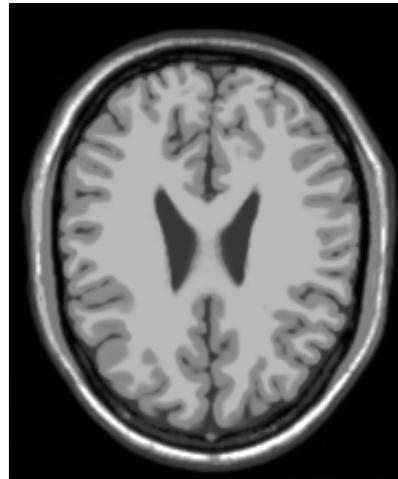
MNI Scaling
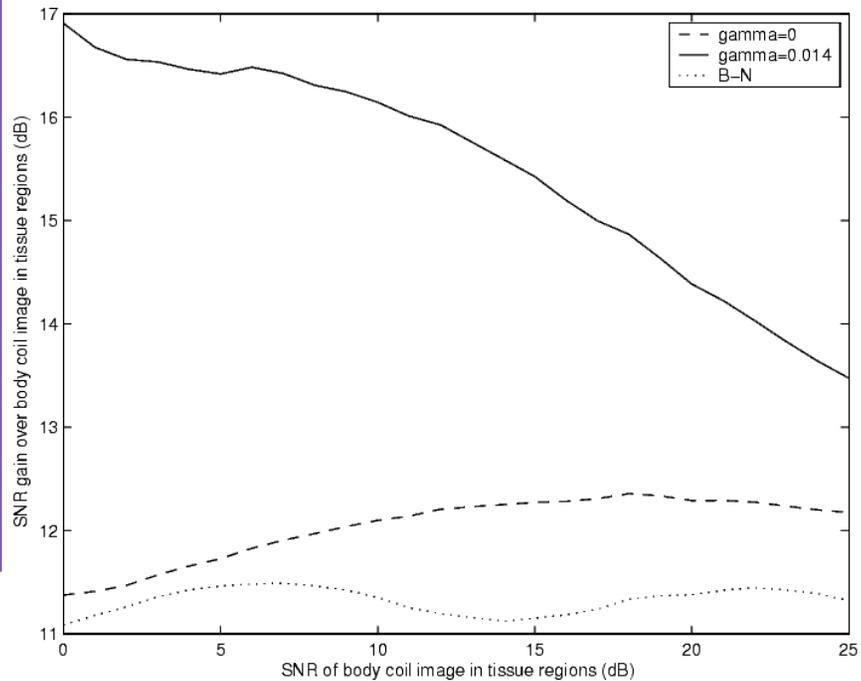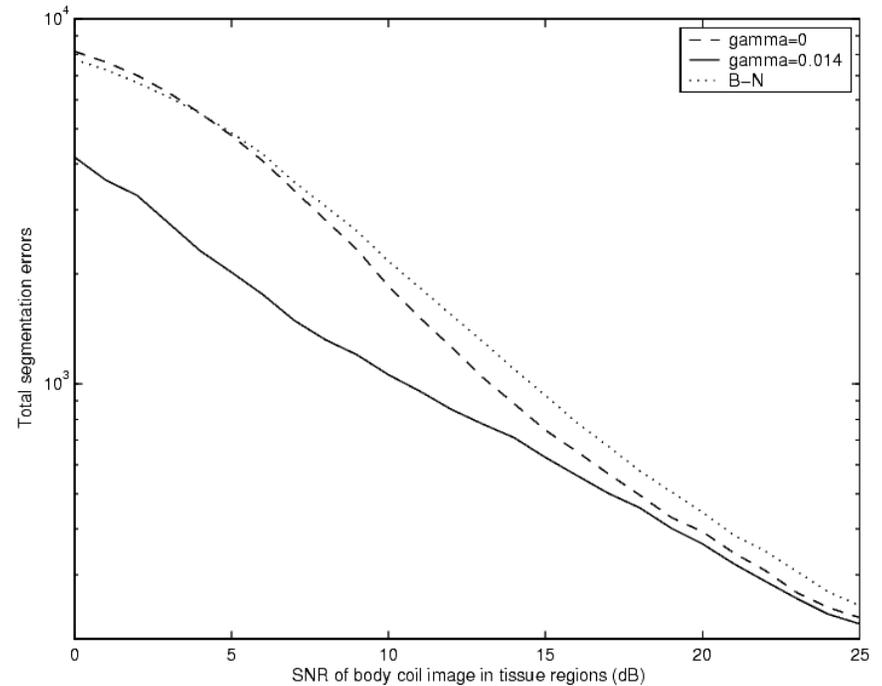
Performance of (a) SNR gain and (b) segmentation errors as a function of image acquisition SNR

Prostate Results
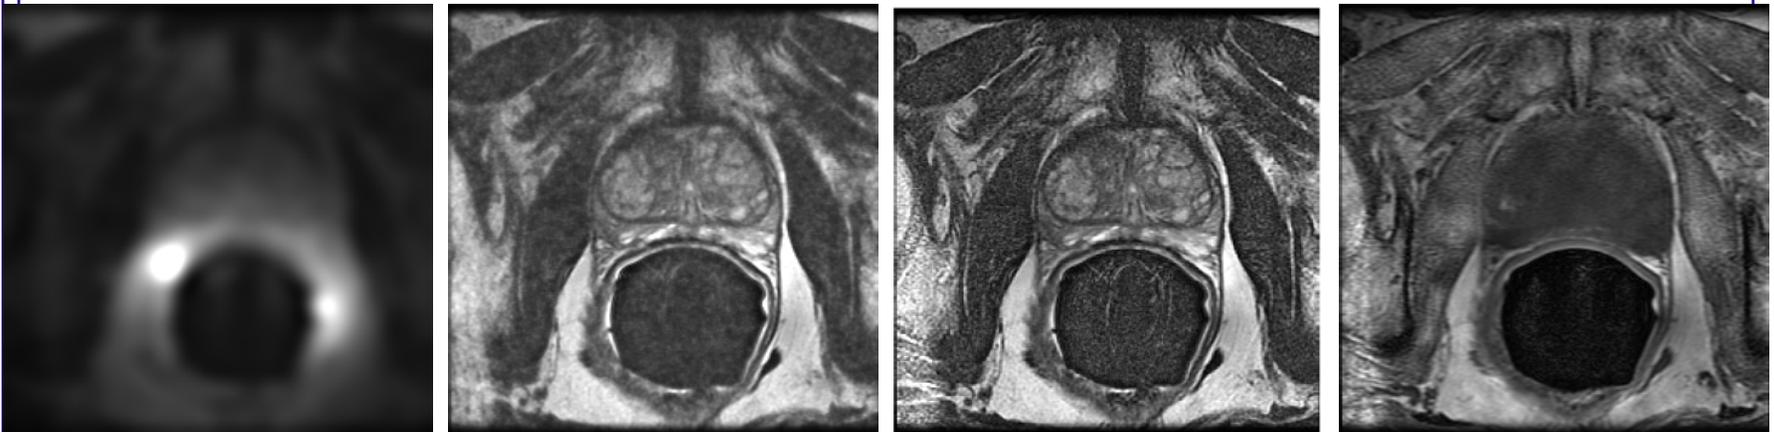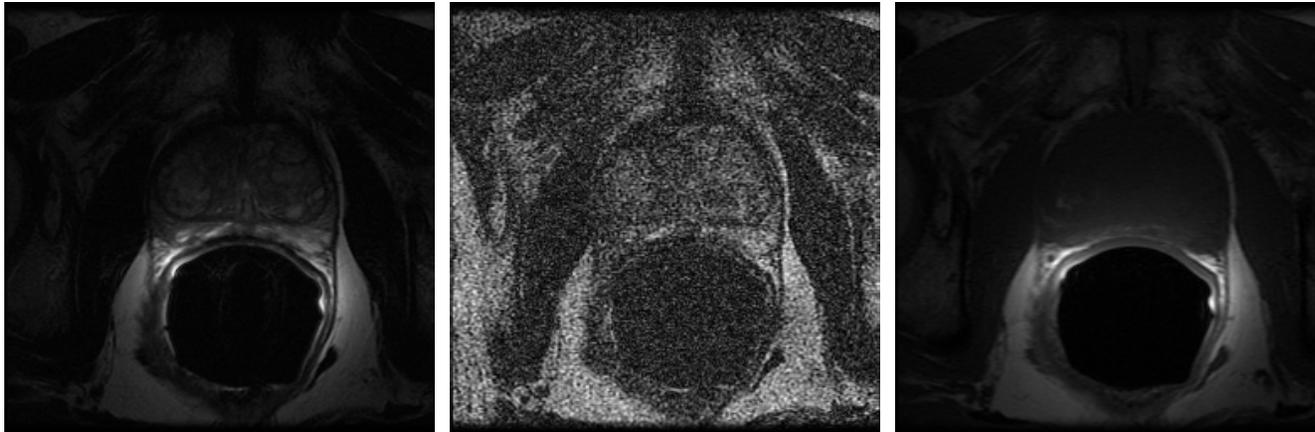
Top: T2W surface coil image, T2W body coil image, T1W surface coil image.
Bottom: Estimated bias field, true T2W image estimate, B-N T2W, true T1W image estimate.

Coronal and Sagittal Correction
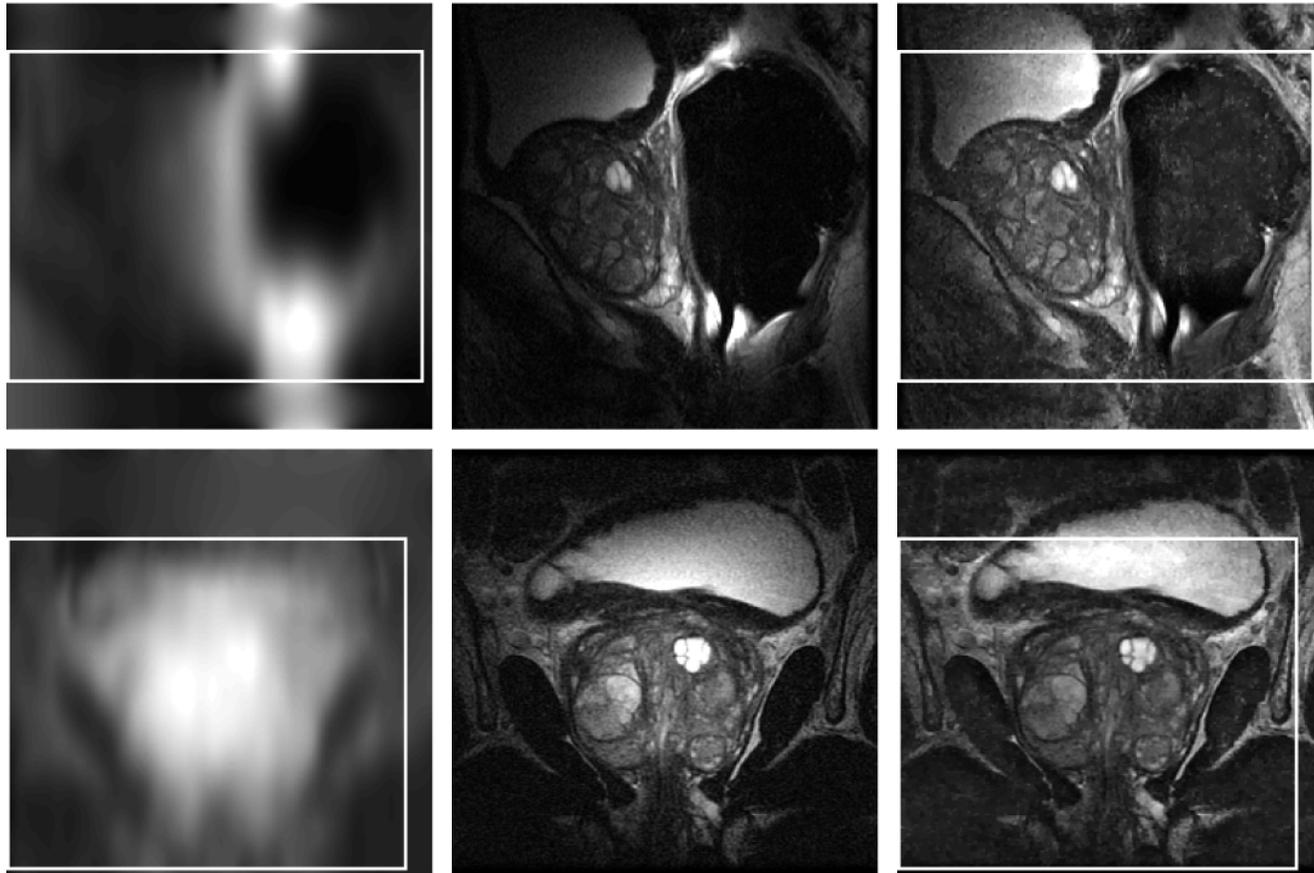
Sagittal (top), coronal (bottom). Bias field (left), surface coil image (middle), true image (right).
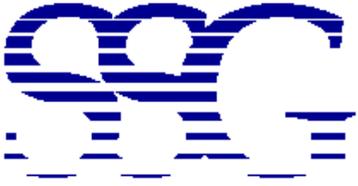
Heart Example
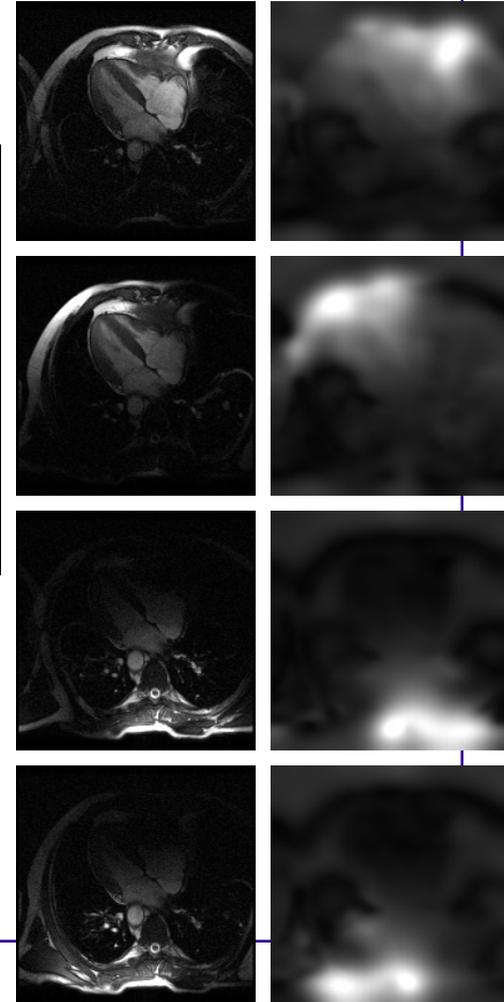
(left to right) Body coil image, B-N estimate, our estimate of f, surface coil images, and bias field estimates.

Heart Movie
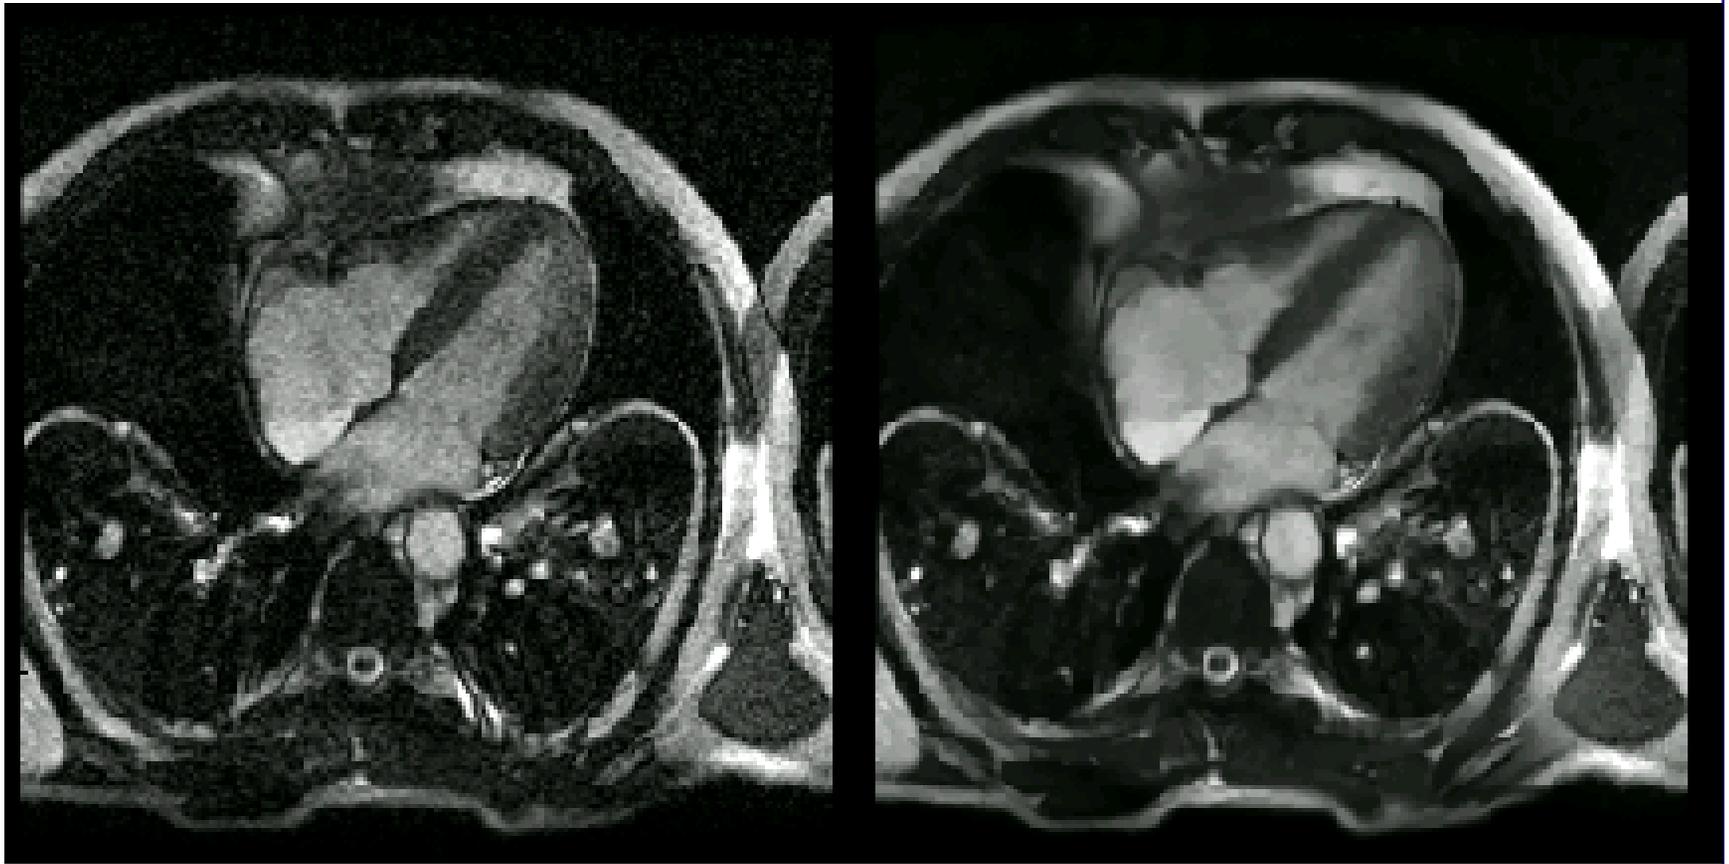
Brain Example
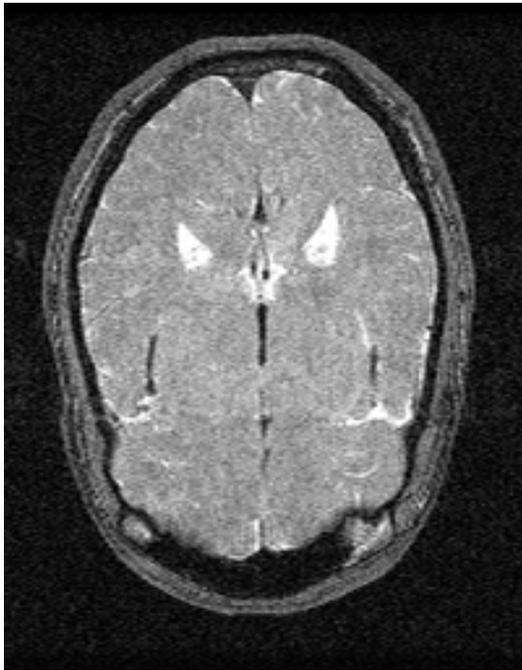
Body coil image (left) and
surface coil images (right),
gradient recalled echo (GRE)
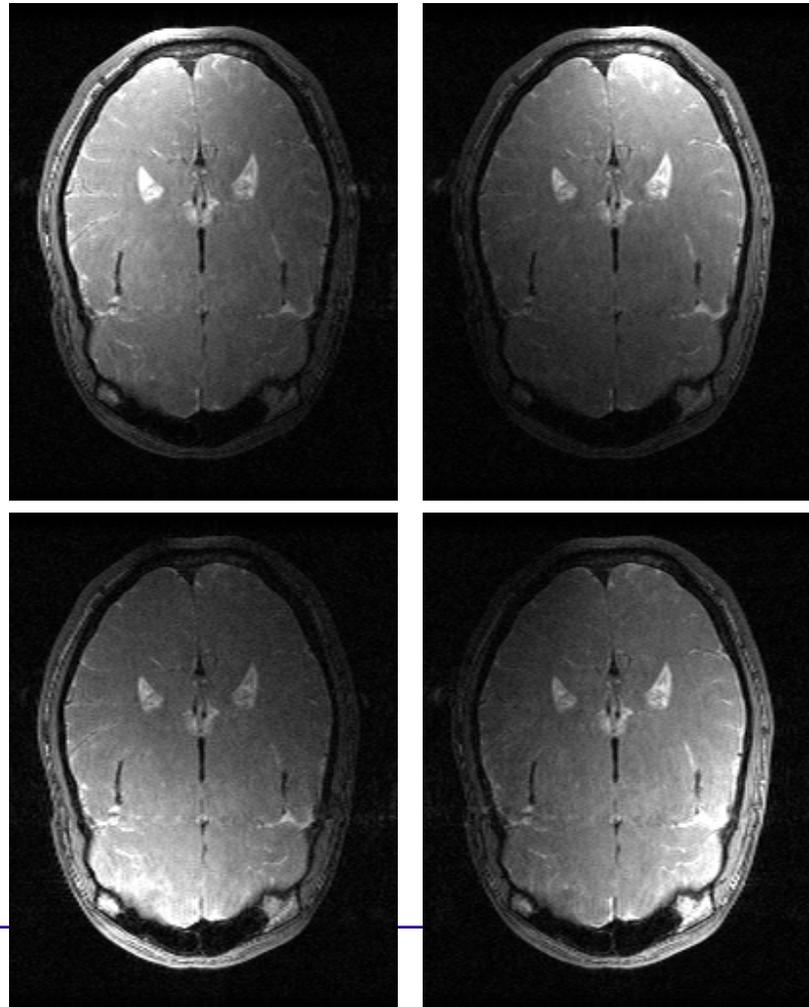
Brain Results
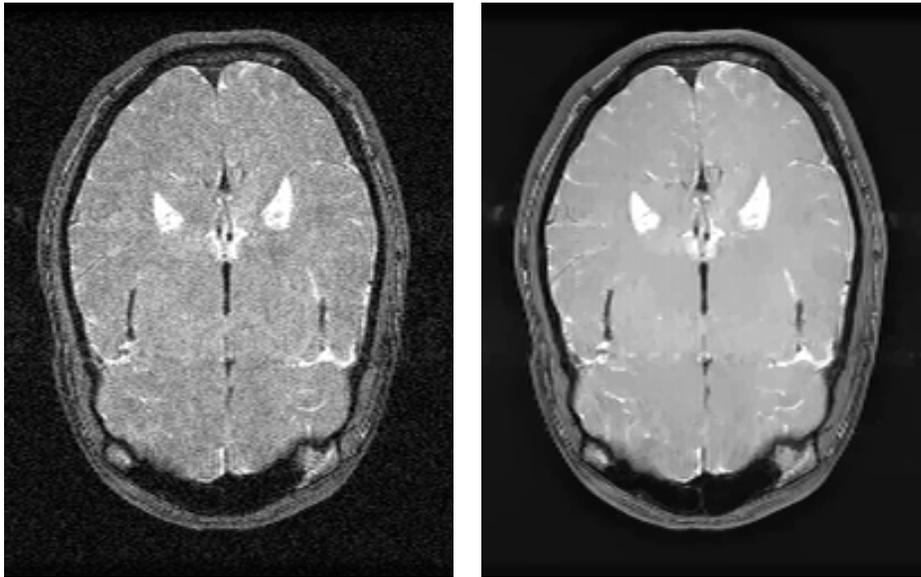
Body coil image (left),
estimated image (middle)
and bias fields (right)
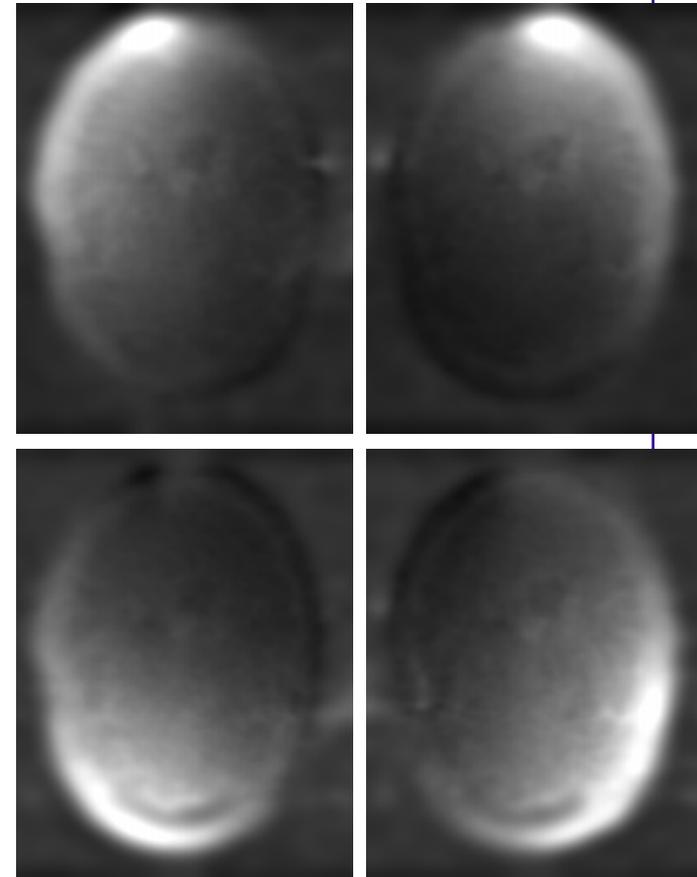
Brain FLAIR Results
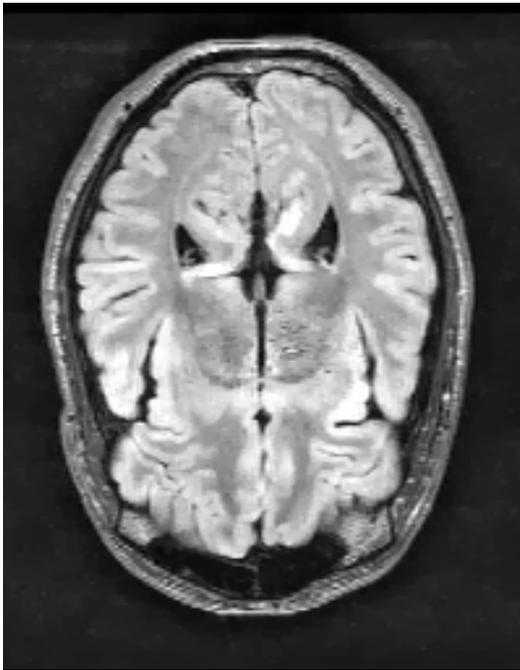
Estimated true image (left),
surface coil images (right),
FLAIR pulse sequence

Outline

Natural Image Model

- We model a natural image as the product of a reflectance map and an illumination map:

$$I_i(x) = r_i(x)l_i(x) + n_i(x), \ i \in \{1, 2, 3\}$$

- The noise is Poisson.

- Both r and l have edges.

- These are color images so we have to deal with each RGB component.

Flash

- We will consider the case of having multiple illuminations of the same scene.

- The specific setup we will use is a picture taken with ambient illumination and a picture taken with flash illumination:

$$I_{a,i}(x) = r_i(x)l_{a,i}(x) + n_{a,i}(x)$$
$$I_{f,i}(x) = r_i(x)l_{f,i}(x) + n_{f,i}(x)$$

Flash Correction

- A simple thing we can consider doing is analogous to what we did in the bias correction problem: fuse the two images together to take advantage of high SNR close to the flash, and the more homogeneous illumination in the ambient light.

- We can make the following relationships and use our previous formulation:

$$b_i(x) = l_{f,i}(x)/l_{a,i}(x)$$
$$f_i(x) = r_i(x)l_{a,i}(x)$$

- Note: this is really the exact same thing we were doing before except we set the bias field of the body coil equal to 1.

- We can use the exact same energy functional as before, except now we also need to use l_p regularization on the b-field as well (this is also imposing a Gaussian noise assumption on our images).

Fusion Results

Clockwise (from upper right): Corrected with black and white map; color correction map; estimated ambient image

Fusion Results (2)

Clockwise (from upper left): Flash image, corrected image, long exposure time image, ambient image

Intrinsic Image Separation

- Simplifying assumption: illumination has constant hue. Thus our model becomes:

$$I_{a,i}(x) = \alpha_i r_i(x) l_a(x) + n_{a,i}(x)$$
$$I_{f,i}(x) = \beta_i r_i(x) l_f(x) + n_{f,i}(x)$$

- Then with N pixels, we have 6N observations, 5N + 6 unknowns.

- This is a valid assumption if most of the light comes from a single source and the amount of reflected light is low.

- Unfortunately it's still ill-posed!

Ill-Posedness

- There is a fundamental problem because we never observe the fields in isolation. So even with enough observations, the problem can remain ill-posed.

- E.g., say we have four fields and four observations:

$$y_1(x) = a_1(x)b_1(x) \qquad y_2(x) = a_2(x)b_1(x)$$
$$y_3(x) = a_1(x)b_2(x) \qquad y_4(x) = a_2(x)b_2(x)$$

- Note that these functions will produce the exact same observations (where c is an arbitrary field):

$$\tilde{a}_1(x) = c(x)a_1(x) \qquad \tilde{a}_2(x) = c(x)a_2(x)$$
$$\tilde{b}_1(x) = b_1(x)/c(x) \qquad \tilde{b}_2(x) = b_2(x)/c(x)$$

Energy Function

- Regularization can be used to make the problem better posed if chosen correctly.

- We construct another energy function with l_2 data fidelity terms and l_p regularization:

$$E(\{r_i\}, l_a, l_f, \alpha, \beta) = \sum_i \left[\lambda_{a,i} ||y_{a,i} - \alpha_i r_i . * l_a||^2 + \lambda_{f,i} ||y_{f,i} - \beta_i r_i . * l_f||^2 \right]$$
$$+ \gamma \sum_i ||Dr_i||_p^p + \eta ||Ll_a||_p^p + \theta ||Ll_f||_p^p$$

Composite Images

Top row: Observed images. Bottom Row: Reconstructed images.

Conclusions

- Non-parametric variational formulation of image fusion problem with statistical estimation flavor.

- Demonstrably superior results on synthetic examples.

- Simultaneous bias correction and denoising.

- Seamless handling of multiple surface coils and multiple pulse sequences.

- Reflectance/illumination separation needs much better regularization because the two fields are too similar.

- Efficient solver using coordinate descent, preconditioned CG, multigrid.

- Need more investigation into finding more effective regularization so that we do not need to use the body coil (body coil image can be viewed as providing a point-wise prior for **f**).

Convergence

- In many ways, our coordinate descent approach can be viewed as being similar to Expectation-Maximization.

  – Can be viewed as an EM implementation if we believe that the regularization terms really are our statistical priors.

  – Have the same convergence properties as EM: every f-step and b-step is guaranteed to decrease the energy.

- In general, we can only hope to find a local minimum.

  – In practice, we have found very robust convergence properties, even without using multigrid.

  – Initialization with random noise will usually find a reasonable result, and the results seem to be the same local minimum.
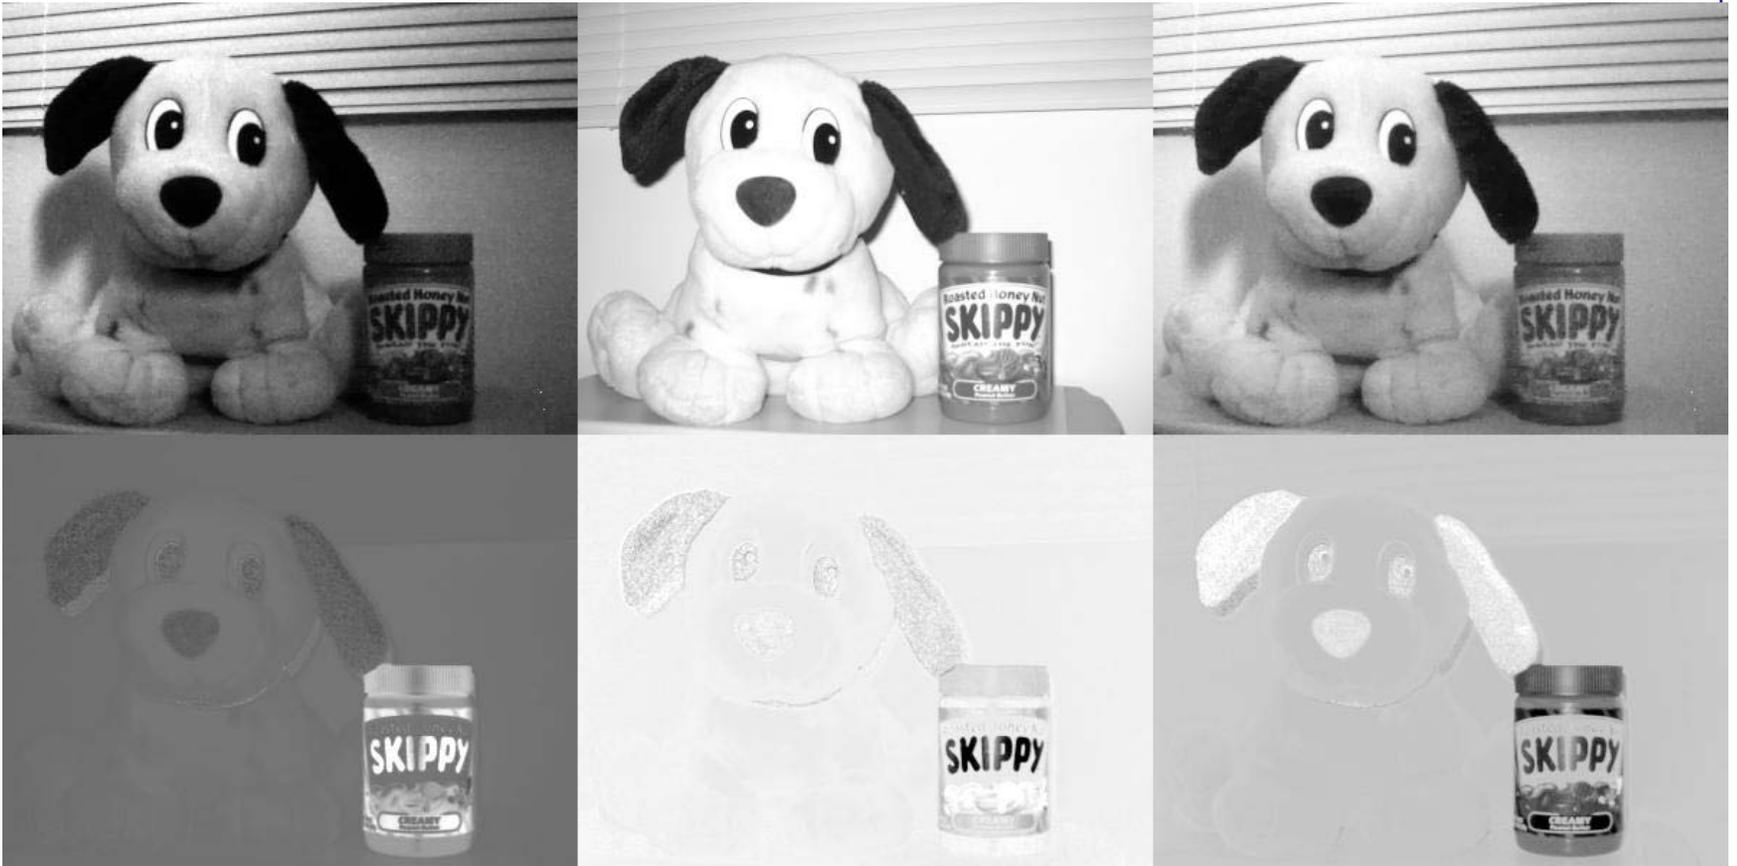
Too Much Dog!

Upper row: illumination maps; Bottom row: reflectance maps

Still Having Problems…

Upper row: illumination maps; Bottom row: reflectance maps