Body-Coil-Constrained Estimation (BoCCE): joint estimation of signal and coil sensitivities from weighted, under-sampled k-space measurements using a novel sampling strategy and associated reconstruction algorithm

M. Dylan Tisdall[1,2]

[1]A. A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, Massachusetts, United States, [2]Radiology, Harvard Medical School, Boston, Massachusetts, United States

Target Audience Developers of parallel image reconstruction algorithms. Users who desire high-quality magnitude and/or phase images with the "flat" sensitivity profile of the body coil, but SNR and undersampling-acceleration gained from a matrix coil.

Purpose We present a novel algorithm for parallel image reconstruction that, when paired with a slightly different acquisition strategy than normally employed in parallel imaging, reconstructs high-quality magnitude and phase images even with substantial undersampling and thus reductions in imaging time.

Method Our novel reconstruction algorithm is related to previous work using the Gauss-Newton method to jointly estimate the coil sensitivities and "true" signal using a smoothness constraint on the coil sensitivities [1]. We previously expanded on this idea by showing that the smoothness constraint could be expressed implicitly by representing the coil sensitivities with compact representations in the Fourier domain (similar to the constraint on the kernels estimated by GRAPPA [2]); a formulation that allows for significant reduction in the number of variables being optimized [3]. However both our previous work and that presented in [1] suffered from an inability to distinguish variations in coil sensitivities from slowly varying changes in "true" signal intensity. In the present work, we add to these ideas an intuition borrowed from how coil sensitivities are commonly estimated for use in SENSE [4]: an image is taken with the body coil, and a second image taken with the matrix coil, with image division giving the sensitivities of each matrix channel on the assumption that the body coil image is "flat". We use this idea to constrain our reconstruction by incorporating k-space samples measured with the body coil in addition to our measurements acquired with the matrix. We then jointly estimate 1) the "true" k-space signal weighted with the body coil sensitivity, and 2) the compact Fourier representation of each matrix channel's sensitivity convolved with the Fourier representation of the inverse of the body coil's sensitivity (*i.e.,* for each channel, the kernel that, convolved with the estimated body-coil-weighted signal, gives the estimate of the k-space for that channel). On the assumption that the body coil's sensitivity is extremely compact in the Fourier domain (*e.g.,* just a DC term), then we expect that the kernels estimated in 2) will be almost as compact as the true channel sensitivities. Performing an inverse FFT on the estimated body-coil k-space estimated in 1) gives an estimate of the "true" image, weighted with the body coil's sensitivity.

The estimation problem is framed as a minimization of the weighted mean squared difference (error) between the complex measured k-space data and the complex k-space data predicted by the estimated parameters (1 and 2 above). In the present example, weights were set to 0 for k-space samples that were not acquired, and 1 for samples that were acquired (separate weights for each channel allow us to distinguish samples acquired with the body coil from samples acquired with the matrix). Our algorithm does not assume binary weights, and so this method could be expanded, at no additional computational cost, to account for additional weighting metrics (*e.g.,* from motion-tracking) at each point in k-space. Minimization was then performed using a Levenberg-Marquardt (LM) algorithm, with a diagonally-preconditioned conjugate gradient algorithm iteratively solving each LM-step [5]. The chosen parameterization of our model allows us to perform these operations on Cartesian-sampled data without forming any of the matrices normally implied by the LM algorithm. Instead, we decompose the "matrix-times-vector" operation into convolutions, element-wise vector multiplications, and inner-products with shifted versions of the vector. These operations are highly parallelizable, and so we developed an initial implementation of our optimization algorithm in C++ and OpenCL [6] so that it could run on a GPU.

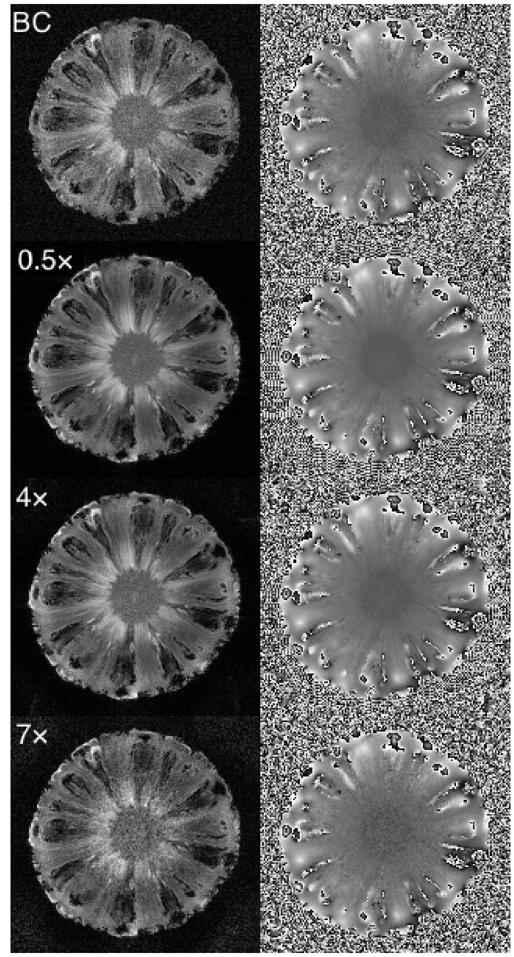

Fig 1. Representative slice from reconstructed volumes. Rows from top: full body coil (BC), full body and matrix (0.5×), 2×2 undersampled (4×), 3×3 undersampled (7×). Columns: magnitude image (left) and phase image (right)

A pineapple was scanned using a 3 T TIM Trio (Siemens Healthcare, Erlangen, Germany) and the product 32-channel head matrix. Two fully-sampled 3D FLASH volumes were acquired sequentially, one with the matrix and one with the body coil. The sequence had a 152×160×256 matrix, $(1\text{ mm})^3$ resolution, 12 ms TR, 3.43 ms TE, 10° flip angle, 200 Hz/px bandwidth, 2× readout oversampling, and 4:52 scan time. We also acquired two noise-only scans, with no pulses, to allow us to estimate the thermal noise covariance matrix of the head matrix and variance of the body coil. Using these, the raw k-space data was linearly transformed so the thermal noise had a standard normal distribution in all channels, and readout oversampling was removed. An image volume was produced from the fully-sampled body coil data as a baseline for comparison. A second volume was produced using our algorithm with fully sampled body coil and head matrix data (0.5× acceleration compared to a single matrix-only acquisition). A third volume was produced with our algorithm where the center 10×10 phase-encode steps were retained from the body coil, while the matrix retained the center 10×10 region and a 2×2 hexagonal undersampling pattern in the rest of k-space, giving 3.89× acceleration compared to a single scan (effective 1:15 scan time). A fourth volume was produced with our algorithm where the center 20×20 phase-encode steps were retained from the body coil, while the matrix retained the center 20×20 region and a 3×3 skewed-hexagonal undersampling pattern in the rest of k-space, giving 7.08× acceleration (effective 0:45 scan time). In all cases, our algorithm used a 7×7×7 representation of the matrix coils and was run on a NVIDIA Tesla C2050 (NVIDIA, Santa Clara, CA, USA).

Results and Discussion Fig. 1 shows a representative slice (readout through-plane, phase/partition encoding in-plane) from the four reconstructed volumes. Reconstruction times for the complete volumes were 41 minutes (0.5×), 58 minutes (4×), and 60 minutes (7×). Although the 7× has a significant SNR penalty in voxels further from the coil elements (expected with reduced scan time), all of the reconstructed volumes are similar to the body coil image, and show no significant artefacts.

Conclusions We have demonstrated a novel reconstruction algorithm for undersampled multichannel data that includes some samples measured with the body coil. The algorithm gives high-quality results even with large acceleration factors. Current runtimes are not practical for clinical use, but further optimization of the software and better GPUs will likely improve this substantially.

References [1] Uecker et al. "Image reconstruction by regularized nonlinear inversion – joint estimation of coil sensitivities and image content" MRM 2008, 60(3):674-682 [2] Griswold et al. "Generalized autocalibrating partially parallel acquisitions (GRAPPA)" MRM 2002, 47(6):1202-1210 [3] Tisdall et al. "Joint Estimation of Signal and Coil Sensitivities with a Bilinear Model" ISMRM Workshop on Parallel MRI 2009 [4] Pruessmann et al. "SENSE: sensitivity encoding for fast MRI" MRM 1999 42(5):952-962 [5] Nocedal et al. "Numerical Optimization" 2006 2nd Ed. Springer [6] Stone et al. "OpenCL: A parallel programming standard for heterogeneous computing systems." Comput Sci Eng 2010, 12(3):66-72.

Acknowledgements This work was supported by: NIH K99HD074649, R01HD071664, P41RR014075.