Acquisition and Reconstruction Effects on Image Quality in Variable-Density Sparse MRI

Dimitris Mitsouras[1], Onur Afacan[2], Robert V Mulkern[3], and Dana H Brooks[4]

[1]*Radiology, BWH/Harvard Medical School, Boston, MA, United States,* [2]*Children's Hospital Boston, MA, United States,* [3]*Childrens' Hospital Boston, MA, United States,* [4]*Northeastern University, Boston, MA, United States*

Introduction: Compressed Sensing (CS) reconstructions impose a sparsity constraint in some domain to recover the correct image. It is important to recognize that the resulting image is affected not only by the choice of k-space trajectory and its parameters (e.g., undersampling rate), but also by the reconstruction method and its own parameters. We performed a systematic joint comparison of spiral variable density (VD) trajectories, that are simple yet known to be ideal candidates for 2D CS MRI (Lustig et al. IEEE Sig Proc Mag 2008;25:78-82), and common CS reconstruction algorithms, carefully quantifying the tradeoffs – such as SNR and reconstruction accuracy – with respect to the selection and parameters of both trajectories and reconstructions, so as to illuminate their relative effects. Results were obtained for each reconstruction method across 16,000 simulations and 3,200 experimental acquisitions.

Methods: Four undersampling schemes were explored: linearly and quadratically decreasing, and two piecewise-constant density schemes (respectively, with undersampling only the outer half or a variable, larger outer portion). Each scheme was used to produce 40 separate trajectories across a range of undersampling factor to achieve a k-space coverage from the Nyquist rate (no undersampling, 22cm FOV, 2.54mm resolution, k_r^{max} reached ≈1.97cm^{-1}), to a most undersampled trajectory (k_r^{max} ≈2.76cm^{-1}, nominal resolution 1.82mm). Five reconstruction methods were explored: two linear, namely direct FFT (Pipe-Menon DCF) and minimum least squares (L2), and, 3 CS, namely minimum L1-norm, Total Variation (TV)-norm, and, combined minimum L1/TV-norm. Results were obtained for a simulated analytic phantom and the ADNI "Magphan" phantom imaged using a GRE acquisition (5mm slice, 30° flip, 200/2.2ms TR/TE; Figure) at 1.5T (GE HDx). Both phantoms were highly sparse in all transform domains used for CS reconstruction(88-62% zero coefficients). Acquisition with each VD scheme and undersampling factor was performed 100 times with independent noise for simulations, and 20 times for experiments. SNR was quantified using the standard deviation (SD) of pixels across acquisitions, and summarized by its average over the entire FOV as well as over the 10% of the FOV exhibiting the highest SD. A map of individual pixel-wise SDs was used to assess noise texture. RMS error was calculated w.r.t ground-truth (analytic simulated and high-resolution high-SNR GRE images for simulations and experiments respectively), and summarized by its mean across the FOV, and over the 10% of the FOV exhibiting the highest residuals. An "error probability map" was defined by the number of times a pixel was in the highest quartile of residuals.

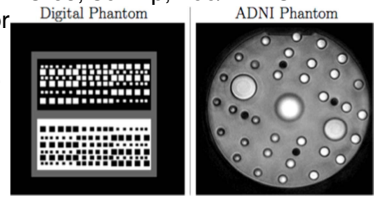

Results: The study required >2 years of CPU time on top-of-the-line processors, and summarized >75GB of 256x256 2D image reconstructions. SNR varied only slightly between trajectories, in close agreement to theory (Tsai et al. MRM 2000;43:452-8). Reconstruction method choice had a much larger effect than trajectory on average image SNR (Fig 1a). Moreover, pixel SD was nearly constant throughout the FOV for linear reconstructions (DCF and L_2) as expected, but with CS reconstructions, the 10% of the image region with highest noise had 2- to 4-fold lower SNR than the average (Figs 1b,2a). These higher SD pixels were clustered in regions mirroring the "inverse" of the transform functional, i.e., those with non-zero signal intensity for L1, and those with signal transitions for TV reconstruction (Fig 2a columns 3,4). Linear reconstructions had higher RMSE than CS (Fig 1c,d). With linear reconstruction, higher undersampling reduced the likelihood of large error at locations of high-resolution features; for example, pixels at the center of small rectangular objects in the phantom affected by Gibbs ringing at the Nyquist rate had reduced error counts at higher undersampling (Fig. 2b). However, high error counts increased in a ripple pattern outside the sample as expected due to aliasing. In contrast, for CS reconstructions, errors outside the sample did not increase at higher undersampling; instead, more pixels of the phantom became afflicted, although less often each. For both SNR and RMSE, the performance of CS reconstruction across the entirety of the FOV was better compared to linear reconstruction and was also directly related to the sparsity of the imaged sample in the transform domain (the more sparse, the better the performance). However, for the 10% of the FOV with highest noise, CS reconstruction had much higher uncertainty (as defined by pixel SD) than linear reconstruction. Experimental results closely followed the simulations.

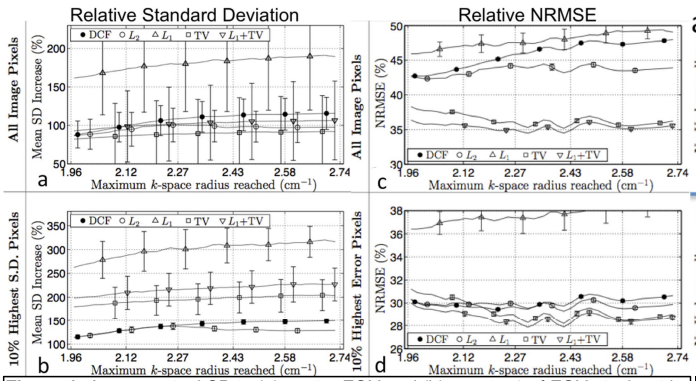

Figure 1. Average pixel SD in (a): entire FOV and (b): top 10% of FOV pixels with highest SD, normalized for an SNR=10. Normalized RMSE in (c): entire FOV and (d): top 10% of FOV pixels with highest RMSE. Images reconstructed using 5 methods (2 linear, 3 CS) over 100 noisy simulations using the quadratic density spiral scheme across 40 undersampling factor choices.

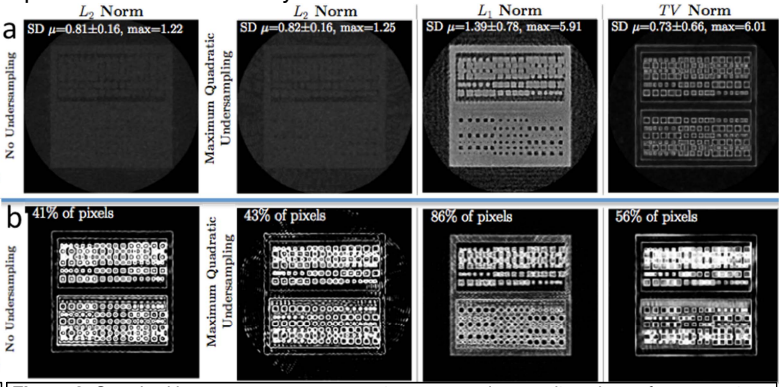

Figure 2. Standard least squares reconstruction at no undersampling shown for comparison in 1st column. (a), top row: SD maps, and (b), bottom row: error probability maps from 100 simulations at maximum quadratic undersampling by 3 reconstruction methods: L2, L1, and TV. Mean and max pixel SD, and % of pixels with non-zero upper quartile error probability are noted in each panel.

Conclusion: CS MRI yields more accurate images than linear reconstruction. However, the higher accuracy is accompanied by a tandem increase in uncertainty (driven by changes in signal noise) as to which transform-domain coefficients are suppressed by the reconstruction algorithm from one acquisition to the next. When the transform domain directly correlates to spatial features, there is increased uncertainty as to whether the final image contains a certain visual feature or not – eg. a hypointensity with L1 reconstruction – due to the particular instantiation of noise. Our results indicate that trajectory choice has a very limited role in this effect; instead, reduced uncertainty likely requires careful assessment and selection of the CS transform domain and reconstruction parameters.