Automatic L_1-SPIRiT Regularization Parameter Selection Using Monte-Carlo SURE

Daniel S Weller[1], Sathish Ramani[1], Jon-Fredrik Nielsen[2], and Jeffrey A Fessler[1]
[1]*EECS, University of Michigan, Ann Arbor, MI, United States,* [2]*BME, University of Michigan, Ann Arbor, MI, United States*

Introduction

L_1-SPIRiT [1] incorporates ℓ_1-norm regularization into the auto-calibrating parallel imaging reconstruction method SPIRiT [2]. Although such regularized methods promise to improve image quality, allowing greater undersampling, selecting an appropriate value for the regularization parameter can impede practical use. This paper proposes a parameter selection criterion based on Stein's unbiased risk estimate (SURE) [3], which estimates the mean-squared error (MSE) without knowledge of the true signal. Traditionally, SURE may be difficult to compute for iterative nonlinear reconstruction algorithms. Recently, a SURE-guided iterative reconstruction [4] was developed for single-channel MRI; we extend this work to multi-channel L1-SPIRiT reconstruction with arbitrary k-space undersampling using a complex-valued Monte-Carlo-based approach. Here, we apply it to reconstruct real parallel-imaging data with varying levels of noise.

Theory

The L_1-SPIRiT estimator $\hat{x} = f_\lambda(y)$ minimizes the wavelet-domain joint sparsity of the multi-channel images from full k-space x that preserves the acquired data y (where the observations have noise covariance Ω) and is consistent with the SPIRiT kernel convolution matrix G; i.e. $x = Gx$. Joint sparsity is enforced using wavelet-domain soft-thresholding with parameter λ; the choice of this parameter determines the sparsity (and the quality) of the reconstruction. We would like to choose λ to minimize the MSE, but since only the acquired data is available, we use SPIRiT to provide the full k-space from the true values at the sample locations. Since the true values at the sample locations are unknown, we turn to SURE to approximate the MSE and form a Monte-Carlo estimate [5] using only a single realization of a complex iid zero-mean unit-variance random vector b and two evaluations of f_λ: $f_\lambda(y)$ and $f_\lambda(y+\varepsilon b)$, where we use $\varepsilon = 10^{-5}$.

Methods

We initially validated this approach using simulated data; here, we show our results for real data. We acquired T1-weighted brain data of a consented subject using 3D SPGR (TR/TE/FA = 25ms/5.2ms/25°) on a 3 T GE scanner with an eight-channel coil, and the noise covariance for this coil was measured from a noise-only pre-scan. We extract a 256x144 axial slice with 1.0×1.25 mm spatial resolution, and we Poisson-disc undersample the data by a factor of 4, retaining a 24x24 block of calibration data. The sparsifying transform for L1-SPIRiT is a four-level \`db4' wavelet; each L1-SPIRiT reconstruction consists of 20 iterations of the POCS implementation [1]. Un-regularized SPIRiT uses 20 CG iterations. The SURE values depicted in Figure 1 contain a constant term that depends on the true sample values; this constant does not affect the minimization and is only included to ease comparison with MSE. For the second experiment, we choose the SURE-optimal λ's via two-level coarse-to-fine parameter sweeps. The SNR of the real data is 16.7 dB; we add complex Gaussian noise to reduce the SNR by factors of 2, 4, 8, and 16. The MSE and signal-to-error ratio (SER) values are computed with respect to the fully-sampled acquired data.

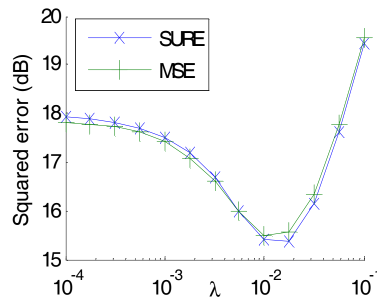

Fig. 1: SURE nearly matches the MSE.

Results

Figure 1 compares the SURE estimate and MSE (using the acquired full k-space as ground-truth) for L1-SPIRiT reconstructions from data with 16x amplified noise for λ's from 10^{-4} to 0.1. These error criteria closely match in shape, with a maximum deviation of 0.2 dB. The second experiment uses SURE-optimal λ selection for a range of noise levels. Figure 2 shows the optimal λ's and reconstructed image SERs for both SPIRiT and L1-SPIRiT (with the SURE-optimal λ) at these noise levels; as SNR decreases, increased regularization leads to greater improvement over un-regularized SPIRiT. Figure 3 depicts sum-of-squares reconstructed images and difference images for the 4 × amplified noise case.

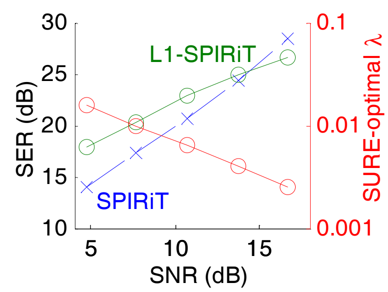

Fig. 2: L1-SPIRiT (green) outperforms SPIRiT (blue), while the SURE-optimal parameter λ (red) is higher, at low SNR.

Discussion

These results using real data suggest that this Monte-Carlo method for automatic regularization parameter selection is effective for L1-SPIRiT reconstruction. We expect this method can be used to tune L1-SPIRiT without user intervention, which is useful in practice. Because Monte-Carlo SURE is a black-box approach that uses only the output of the reconstruction and not the structure of f_λ, this implementation can be generalized to other parallel imaging reconstruction and denoising methods without substantial adjustment.

References

1. Murphy M, Alley M, Demmel J, Keutzer K, Vasanawala S, Lustig M. IEEE Trans Med Imaging. 2012;31(6):1250-62.
2. Lustig M and Pauly JM. Magn Reson Med. 2010;64(2):457-71.
3. Stein C. Ann Stat. 1981;9(6):1135-51.
4. Ramani S, Liu Z, Rosen J, Nielsen J-F, Fessler JA. IEEE Trans Im Proc. 2012;21(8):3659-72.
5. Ramani S, Blu T, Unser M. IEEE Trans Im Proc. 2008;17(9):1540-54.

Acknowledgments

NIH F32 EB015914 and NIH/NCI P01 CA87634.

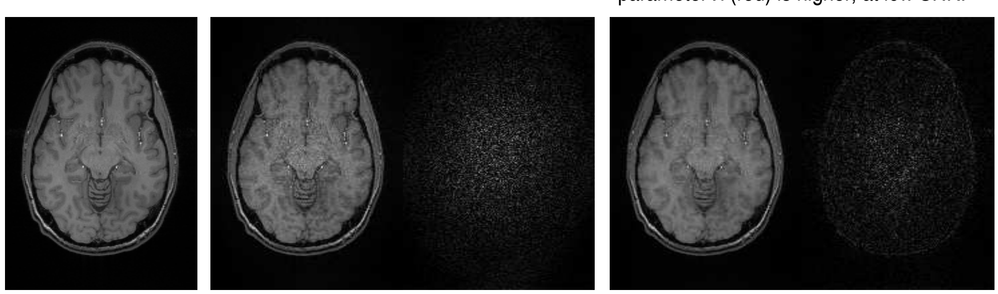

Fig. 3: Reconstructed and difference images for SPIRiT (middle) portray significant noise amplification, while the noise level in the L1-SPIRiT (right) output with SURE-optimal λ is significantly reduced, almost matching the original image (left). The difference images are windowed up by a factor of 4.