Compressed Sensing with Self-Validation

Yudong Zhu[1]
[1]*Zhu Consulting, Scarsdale, NY, United States*

Introduction: Compressed sensing holds much potential for advancing clinical MR[1]. Examples abound demonstrating its success in accelerating MR data acquisition. Theoretical analyses offer not only insights on its statistical behavior but asymptotic performance guarantees under certain conditions[2,3]. In reading an image obtained in a specific imaging instance however, it is not uncommon for one to have a lingering concern that the underlying imaging scheme might have, in a convoluted and subtle manner, obscured some diagnostically important features. Compared to traditional imaging schemes that rely principally on linear operators, a compressed sensing scheme, with at its core a nonlinear operator that leverages random sampling and a sparse model to untangle signals from interferences, poses more challenges to gauging the level of image fidelity. This is because the scheme's response to signal, noise and reconstruction parameters is difficult to grasp or interpret. In this work we investigated the feasibility of introducing self-validation into compressed sensing MR. The goal is to assist image fidelity assessment/improvement in practice with validation tests that can be automatically performed on any specific imaging instance itself, without requiring additional data or comparison references.

Method and Results: Post-scan tests were devised to detect issues that are difficult to spot by inspecting alone, image(s) in the instance. The foundation for the tests include: a) Critical reliance on a nonlinear operator in reconstruction calls for at least a validation of the operator's response at or near the operating point (as set by the acquired data). b) Small perturbations to an imaging instance should not significantly impact its outcome if the instance is intrinsically robust or reliable. c) Linear response to image features is essential to contrast fidelity. d) From the perspective of compressed sensing reconstruction, the principle of random sampling implies that equivalence exist among different k-space sampling patterns (e.g., two sets of k-locations both selected uniformly at random[2] can yield identical reconstruction results).

Linear Response Test (L-test) An indication of an imaging scheme's inability to linearly track a local signal change shall call into question the accuracy of image contrast/feature at that location. To capture such indications an L-test checks the difference between an image in the instance and an image reconstructed from perturbed k-space data that correspond to adding a local "bump" to the true image. For example the test checks the difference between a pair of results from, respectively:

$$\text{argmin}_m \quad \|\mathcal{F}_u m - y\|_2^2 + \lambda\|\Psi m\|_1 + \alpha TV(m) \qquad \text{and} \qquad \text{argmin}_m \quad \left\|\mathcal{F}_u m - (y + \mathcal{F}_u g^{(i)})\right\|_2^2 + \lambda\|\Psi m\|_1 + \alpha TV(m) \qquad [1]$$

where $g^{(i)}$ represents a "bump" centered at pixel i – in this study $g^{(i)}$ represents a 5x5 Hamming window profile centered at pixel *i*. Through $\mathcal{F}_u$, the under-sampled Fourier transform[1], the perturbation is inflicted broadly in k-space. At pixel *i*, deviation of the value of the difference from that of the "bump" indicates nonlinearity in local signal recovery. Sweeping a region and recording the deviation pixel-by-pixel creates a deviation-from-linearity map that flags contrast fidelity issue in the region. Ideally the reconstruction outcome responds linearly to the measured data *y*, which preserves contrast and ensues zero deviation. The sampling pattern, sparse model and numerical solver involved in a specific imaging instance however may situate the reconstruction in a nonlinear regime, causing deviation detectable by an L-test.

Sampling Test (S-test) An S-test checks outcome variation amongst an ensemble of compressed sensing reconstructions that each uses a leaving-one-out version of the original set of acquired k-space data. For example the S-test checks standard deviation of the solutions to the following ensemble of problems:

$$\text{argmin}_m \quad \left\|\mathcal{F}_u^{(j)} m - y^{(j)}\right\|_2^2 + \lambda\|\Psi m\|_1 + \alpha TV(m) \quad , \quad j = 1, 2, \ldots \qquad [2]$$

where $\mathcal{F}_u^{(j)}$ and $y^{(j)}$ represent the *j*th instance of leaving one k-space sample out. Conceptually, if *M* k-space samples are more than adequate for a full reconstruction[2], randomly leaving one sample out without altering the original sampling pdf may not impact the reconstruction. Along this line of reasoning a significant variation detected by an S-test is an indication that the imaging instance's result is not reliable. From a more general perspective, excessive sensitivity of a reconstruction's outcome to perturbation is never a good sign. An S-test involves local perturbations to k-space.

We appreciate Dr. Lustig's online publication of *sparseMRI* [4] and support for reproducible research. The data and code contained therein were used to set up simulation studies in this work.

A first study simulated acceleration of 2D phase encoding in imaging a 256x256 brain slice. The k-space was sampled with patterns that correspond to various acceleration factors, fully-sampled lower k-pace radii and probability density functions – Case 1: 2x and *r*=0.32, Case 2: 4x and *r*=0.32, Case 3: 2x and *r*=0, and Case 4: 4x and *r*=0.02. Fig.1 col.1 shows the compressed sensing results. Fig.1 col.2 displays, over the tested region (marked), the deviation-from-linearity maps from performing L-tests pixel-by-pixel over the region (Eqn.1). Notice the near zero deviation in Case 1, modestly elevated deviation in Case 2 (higher acceleration) and Case 3 (smaller fully-sampled lower k-pace), and significant deviation in Case 4 (both higher acceleration and smaller fully-sampled lower k-pace). For each of the cases, the S-test involved a 500-instance ensemble, where each instance used all k-space samples of the case except one sample that was randomly selected outside the fully-sampled lower k-pace. Fig.1 col.3 summarizes the S-test results in the form of standard deviation maps. Having access to the true image – a luxury unavailable in actual practice – allows direct quantification of the difference between true and reconstructed images, as Fig.1 col.4 shows. Intriguingly, the variations uncovered by the tests (col.3) were rather revealing of the actual error (col.4). In Cases 1-4, the rms of the actual error were 2%, 3%, 4% and 7% of peak image intensity respectively, and the rms of the stddev values were 0.01%, 0.02%, 0.04% and 0.06% of peak image intensity respectively. All plots of Fig.1 used a common gray scale. For each of the cases the lead principal component obtained from SVD of the ensemble of instances (not shown) also appeared to be indicative of the case's actual error pattern.

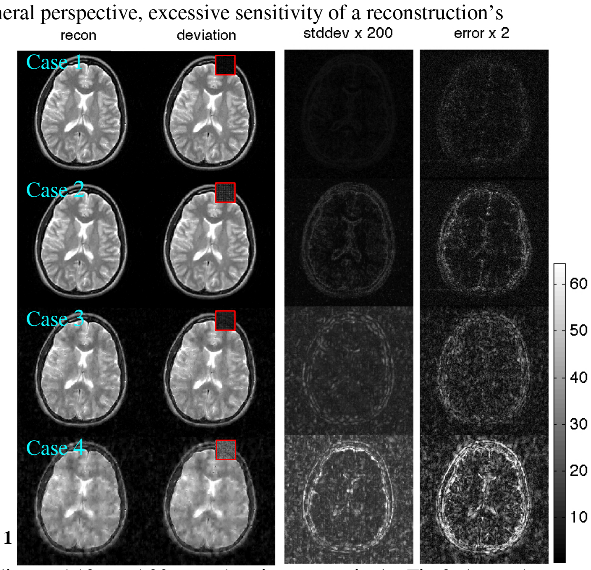

FIG 1

A second study evaluated two angiogram simulations of Ref.1 that were with variable density sampling and 12x and 20x acceleration respectively. Fig.2 shows the compressed sensing results (left column, also in Ref.1), the stddev maps from S-tests (mid column), and the actual error (right column). Notice that the S-tests showed promise in tracking the actual error in this study too. The loss of two low-contrast features in the 20x case for example, was detected by the stddev map.

Discussions: Compressed sensing offers a capacity for speeding up data acquisition while keeping aliasing and noise effects subdued. Theories and experiences however are yet to establish a more robust guidance on random sampling, sparse model and non-linear solver, to help manage the challenge of using the technology in diagnostic MR. The present method addresses the challenge from an alternate angle. When an accelerated scan is done but additional data for validation/comparison is unavailable, the present method steps in to derive fidelity indicators from the data and the reconstruction.
As many imaging examples have suggested, the ingredients of compressed sensing can accommodate linear response and contrast fidelity for a range of setups. A straightforward L-test helps detect if a specific setup is over the range or the resulting image is unreliable. The intrigues of compressed sensing appear to support unique mechanisms for further self-validations. A notion that more (k-space samples) is not necessarily better – think of the challenge a regular under-sampling poses – inspired the idea of randomly excluding one or a few data samples from a specific reconstruction instance. An S-test implements the idea, aiming to detect if and where flaws exist in an otherwise clean looking image. Given the task of imaging a specific, unknown object, no sampling pattern instance is likely to guarantee maximum effectiveness recovering signals from interferences. We hypothesize that an ensemble of compressed sensing reconstructions, each using a random leaving-one-out version of an original set of acquired k-space data, collectively can have a good chance of detecting and/or recovering features that happen to be obscured by a reconstruction using the original set.

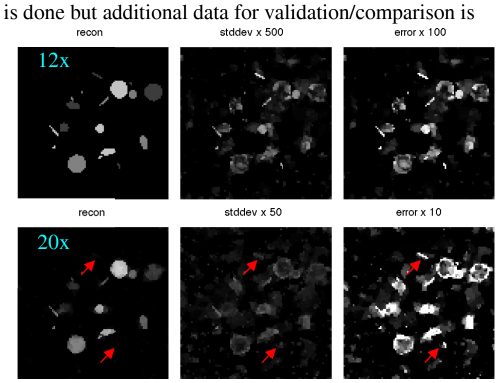

1.Lustig, et al *Magn Reson Med* 58:1182–95, 2007. 2.Candes, et al *IEEE Trans Inf Theory* 52:489–509, 2006. 3.Donoho, *IEEE Trans Inf Theory* 52:1289–1306, 2006. 4. http://www.eecs.berkeley.edu/~mlustig/software/