Communications in Applied Mathematics and Computational Science

SIMULATING SINGLE-COIL MRI FROM THE RESPONSES OF MULTIPLE COILS

MARK TYGERT AND JURE ZBONTAR

vol. 15 no. 2 2020

msp

COMM. APP. MATH. AND COMP. SCI.
Vol. 15, No. 2, 2020
https://doi.org/10.2140/camcos.2020.15.115

msp

SIMULATING SINGLE-COIL MRI FROM THE RESPONSES OF MULTIPLE COILS

MARK TYGERT AND JURE ZBONTAR

We convert the information-rich measurements of parallel and phased-array MRI into noisier data that a corresponding single-coil scanner could have taken. Specifically, we replace the responses from multiple receivers with a linear combination that emulates the response from only a single, aggregate receiver, replete with the low signal-to-noise ratio and phase problems of any single one of the original receivers (combining several receivers is necessary, however, since the original receivers usually have limited spatial sensitivity). This enables experimentation in the simpler context of a single-coil scanner prior to development of algorithms for the full complexity of multiple receiver coils.

1. Introduction

Typical modern magnetic resonance imaging (MRI) scanners (such as all those represented in the fastMRI data of [13]) include *multiple* receivers — so-called "coils" of conductors — for making measurements. In this brief yet self-contained technical note, we seek to conveniently simulate the response of a single-coil machine, complete with errors that are reasonably representative of the measurement errors which a real scanner would make, given only the responses from a multicoil, "parallel-imaging" machine. The fastMRI competition issues two challenges — single-coil and multicoil — under the expectation that the single-coil track could provide a simplified stepping stone toward the full multicoil challenge.

The aim of fastMRI is compressed sensing — the accelerated acquisition of image reconstructions by taking fewer measurements than classical signal processing would need to attain the same resolution; compressed sensing yields full-resolution or "superresolved" reconstructions based on so-called "undersampled" measurements, measurements acquired at lower than the Nyquist rate required to reconstruct any arbitrary signal for a given bandwidth. Two approaches to compressed sensing are especially popular: (1) optimization for individual images and (2) machine learning over collections of images. Combinations of the two approaches are also possible.

MSC2010: primary 68U10, 78A50, 78M50, 92C55; secondary 65K10, 78A55, 94A08.
Keywords: magnetic resonance imaging, parallel imaging, multicoil, fastMRI, deep learning, machine learning.

Techniques representative of the first approach (via optimization) include Fourier reconstruction with total-variation regularizers (or related regularizers, such as the sum of the absolute values of wavelet coefficients); representative of the second approach (via machine learning) is statistical modeling with deep neural networks. The terminology "compressed sensing", "compressive sensing", and "compressive sampling" is due to [1; 2]; they took the first approach (via optimization). The second approach (machine learning) is the focus of [13].

Complicating the development of compressed sensing for MRI is the longstanding focus on leveraging multiple receiver coils in the design of MRI scanners, as reviewed by [5]. All scanners of [13] are multicoil by default, designed and optimized for multicoil measurements. However, compressed sensing is more complex with multicoil data. Distilling the multicoil challenge into a single-coil stepping stone could lower barriers to entry into research on compressed sensing for MRI and accelerate the development of accelerated MRI acquisitions, by teasing apart the fundamental mathematical issues from the multicoil complications.

Thus, we need to simulate a single-coil scanner with sufficient fidelity to actual errors in measurement, using only multicoil data. As hoped (and illustrated with numerical examples in Section 3 below), our single-coil emulation does indeed yield phase similar to the notorious phase problems tackled by [4; 6] and others. Encouragingly, the method of [8] is similar to ours and yielded highly successful results in a related setting (though, admittedly, our mathematical motivation, practical application, optimization procedure and targets, and detailed objective differ). Section 2 now introduces our simulation of a single-coil scanner.

2. Emulated single-coil

We propose using a linear combination of the responses from multiple coils for the emulated single-coil (ESC) response. We least-squares fit the complex-valued coefficients in the linear combination to the "ground-truth" reconstruction, estimating the ground truth using the canonical full multicoil reconstruction, the root-sum-square (RSS) of [7]. Most notably, linearly combining the raw coil responses sums complex-valued fluxes directly, rather than summing nonnegative energies as in the RSS. This simulates a single-coil scanner for the reasons discussed next.

2.1. ***Review.*** MRI scanners originally took measurements using only a single coil of wires, recording the magnetic flux through that coil, separately for each pixel or voxel in the spatial Fourier domain ("k-space"). Most large modern scanners instead take most measurements using several separate coils; the simplest such phased array of coils partitions what would have been one large coil into multiple disjoint smaller coils, as illustrated in Figure 1. Squaring the RSS of [7] sums the squares of the absolute values of the spatial inverse Fourier transforms of the magnetic

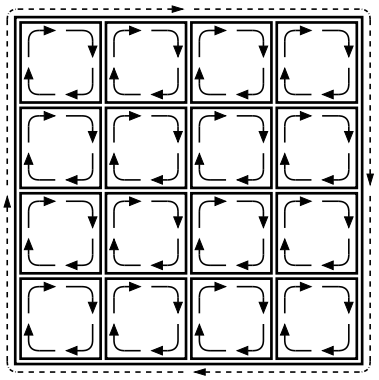

Figure 1. The wide solid lines represent coils of conductors — 16 small ones and 1 large one. The arrows indicate circulation of current through the conductors (with solid lines for the small coils and dashed lines for the large coil). The currents in the small coils would sum to induce the indicated current in the large coil (flux through the small coils would be equivalent to flux through the large coil), neglecting mutual inductances — and actual designs of coil arrays tend to minimize mutual inductances (as per [10]).

fluxes through the coils to obtain the total energy (or power) through the array; the total is the same as what a single coil around the perimeter of the array would have measured for the total energy (or power), but with improved signal-to-noise. (The disjoint partition ensures that the squares of the absolute values of the spatial sensitivities of the coils sum to 1 at any point in space, assuming that the array of coils captures all radiation emanating from the object being imaged.) Most MRI machines use several different partitions of what would have been a single large coil, with those different partitions overlapping each other, again to improve the signal-to-noise ratio. The partitions can leverage polarization to make more orthogonal measurements via quadrature pairs, butterflies, and figure eights; the survey of [10] discusses this together with "bird-cage" and other cylindrical or volumetric coils. In all cases, the basic principle is the same, and the RSS works the same.

2.2. ***Motivation.*** As seen in Figure 1, the flux through a single coil is physically the sum of the fluxes through multiple coils forming a disjoint partition of the single coil that loops around the multiple coils (provided that the multiple coils are all in phase and have the same gain, while neglecting the mutual inductance that coil design typically tries to minimize, as per [10]). This is much like in the proof of Stokes' theorem (which also uses Figure 1). Since we do not know the correct phase offset and gain calibrations, which amount to multiplying a whole coil's measurements by the same complex number, we can fit the complex number so as to match the ground-truth RSS reconstruction as well as possible.

As described, for example, by [10], many MRI machines use "bird-cage" or other volumetric coils rather than surface coils in a flat plane. A phased "single-coil"

array is often a cylindrical array of coils with capacitative (or inductive) couplers between the coils. These couplers effectively introduce phase offsets between the multiple conductors forming the "single-coil" bird-cage. Our linear combination of coil responses accounts for phase offsets and is physically realizable and similar to bird cages used in practice. The use of coils selecting for circular polarization (including the quadrature pairs, butterflies, figure eights, etc., described by [10]) is an additional complication, implicitly resolved by how the least-squares fit of the linear combination to the ground truth decides to combine the coils. The following subsection formalizes the mathematical details.

2.3. ***Mathematical formulation.*** To set notation, we define k to be the number of coils, l to be the number of slices in a volume of cross-sectional images, m to be the height of each slice, and n to be the width of each slice. The measured data from multiple coils then consists of an $(nml) \times k$ matrix A with complex-valued entries. Each of the k coils yields l cross-sectional slices, with each slice being an $m \times n$ image. More precisely, the images in A are cropped inverse Fourier transforms of the original two-dimensional Fourier domain ("k-space") measurements, cropped to the center $m \times n$ block of pixels (the cropping eliminates the usual oversampling to focus on the region of interest to radiologists).

To guide the fitting of parameters for the simulations, we provide "ground-truth" targets by filling an $(nml) \times 1$ column vector b with the RSS reconstructions of [7] obtained from the full multicoil data; specifically, each entry of b is the Euclidean norm of the corresponding $1 \times k$ row in A. While this "ground truth" may be imperfect, the RSS is the canonical choice, used universally throughout medical literature since the work of [7]. Each pixel of each cross-sectional slice in b is the square root of the sum of the energy (or power) flowing for that pixel through the individual coils; the individual coils could be arranged, for example, as illustrated in Figure 1. Each pixel of this RSS reconstruction b collects together all the energy (or power) for that pixel from the individual coils.

Given A and b, we calculate the coefficients in a linear combination of the coils as the (complex-valued) entries in a $k \times 1$ column vector x minimizing

$$\left\| \sqrt{|Ax|} - \sqrt{|b|} \right\|_2^2, \tag{1}$$

where $\| \cdot \|_2$ is the Euclidean norm, $\sqrt{\cdot}$ takes the square root entrywise, and $|\cdot|$ takes the absolute value entrywise, so that $|b|$ is in fact equal to b, as the entries of b are nonnegative. The distance $\left\| \sqrt{|Ax|} - \sqrt{|b|} \right\|_2$ is known as the Hellinger metric of [3]; the square roots amplify entries of small magnitude and attenuate entries of large magnitude. Furthermore, $\sqrt{|Ax|} - \sqrt{|b|} = (|Ax| - |b|)/(\sqrt{|Ax|} + \sqrt{|b|})$, so the Hellinger distance is more like a measure of relative errors than the conventional Euclidean distance $\| |Ax| - |b| \|_2$. Our estimate of $|b|$ (which equals b) is $|Ax|$,

where x minimizes (1). We use absolute values in our estimate $|Ax|$ since the entries of the ground truth we want to match are nonnegative. We use the Hellinger distance since human visual perception is somewhat more calibrated for relative rather than absolute errors; in many cases, omitting the square roots — thereby optimizing mainly for the pixels of greatest intensity — would spoil all but a small fraction. That said, minimizing (1) is far from the only possibility; minimizing (1) performed well in our numerical experiments and facilitates the nonlinear optimization required.

We used the LBFGS (limited-memory Broyden–Fletcher–Goldfarb–Shanno) version of quasi-Newton iterations as implemented in SciPy [11] to solve the nonlinear least-squares problem minimizing (1). We started the iterations with the solution to the linear least-squares problem minimizing

$$\|Ax - b\|_2^2; \tag{2}$$

the objective function in (1) is the same as in (2) aside from taking the square roots of the absolute values of the entries of Ax and of b. As indicated above, the solution x minimizing (2) is unsuitable for the final results, as some entries of Ax can be negative, and only the largest entries of Ax will be accurate in comparison with those optimizing the Hellinger metric. Our final results always use $|Ax|$ with the solution x minimizing (1).

Remark. Anuroop Sriram points out that the mean of the noise on the "ground-truth" RSS will be strictly positive in regions of the image where the actual object being imaged is absent, as the RSS is always nonnegative. Thus, when we fit the linear combination of the coil responses, we should try to match the RSS only where the RSS is substantially above a threshold that is significantly above the noise level. For the current data from fastMRI of [13], essentially the entire cropped image of interest meets this criterion, so we try to match the full cropped image.

3. Numerical examples

We compare the emulated single-coil (ESC) method of Section 2 with the leading singular vector of the singular value decomposition (SVD) of A from (1), as in the "eigencoil" or "coil compression" reviewed by [5]. ESC could be regarded as a special kind of the "virtual-coil" methods discussed by [5]. We also display the RSS of [7] for reference as our best estimate of the ground truth (the second paragraph of Section 2.3 above reviews the RSS).

For our data — scans of 1594 knees from [13] — the parameters from Section 2 take the values $k = 15$ receiver coils, $m = 320$ rows, and $n = 320$ columns in a cross-sectional image, with Figure 2 specifying l, the number of cross-sectional slices. We used a memory of 10 vectors in the limited-memory BFGS quasi-Newton iterations to minimize (1) starting from the linear least-squares solution minimizing (2).

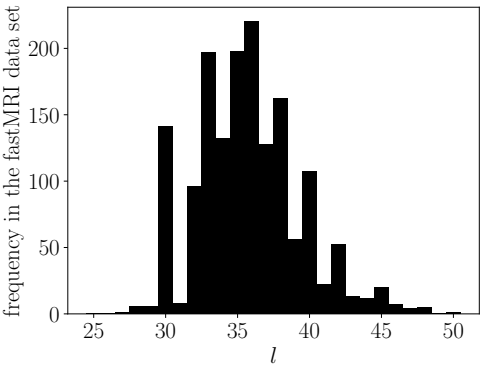

Figure 2. Number of examples in the fastMRI data set of [13] for each possible value of l (the parameter l is the number of cross-sectional slices).

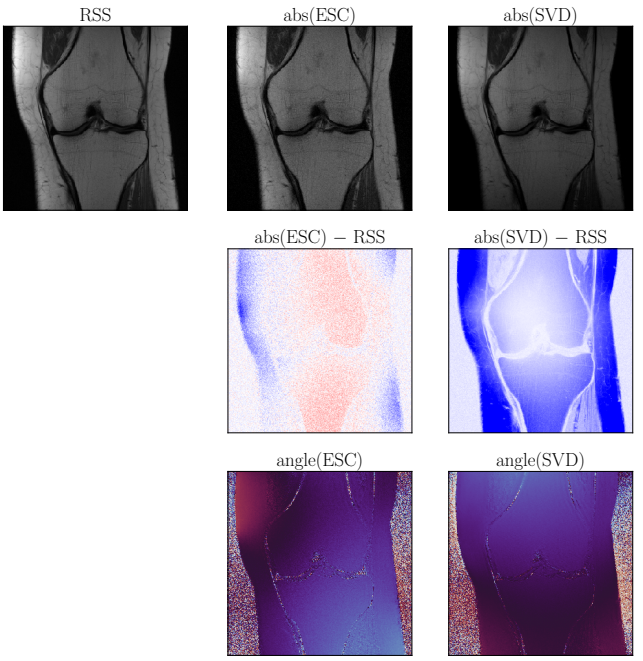

Figure 3. First randomly chosen example; "abs" refers to the absolute value of a complex number and "angle" refers to the phase.

In Figures 3–8, the first row shows (the magnitude of) the reconstructions. The second row shows the difference between the RSS ground truth and the magnitude of each method's reconstruction. The third row shows the phase of the reconstructions (estimates of the RSS ground truth discard the phase). Figures 12–14 in the Appendix display the entries of the solution x minimizing (1), where the ESC reconstruction is Ax.

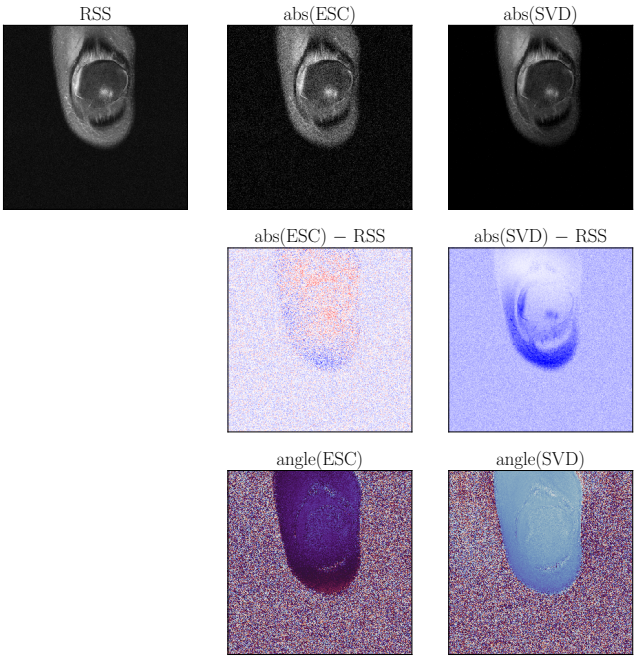

Figure 4. Second randomly chosen example.

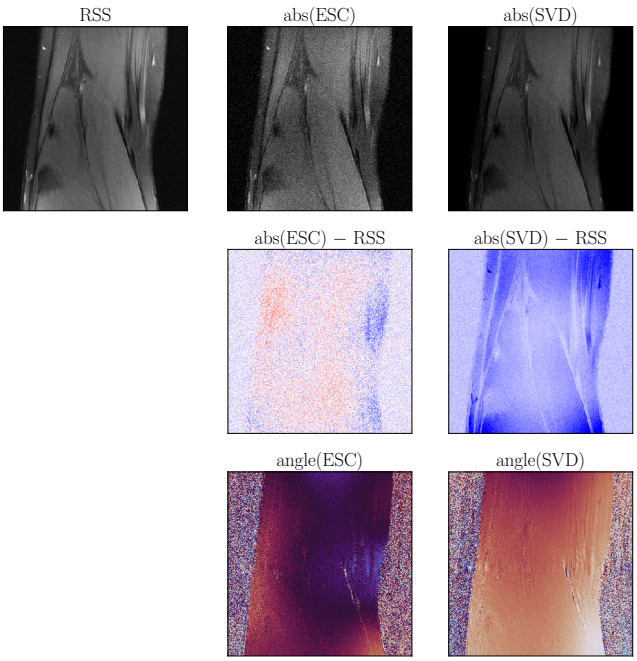

Figure 5. Third randomly chosen example.

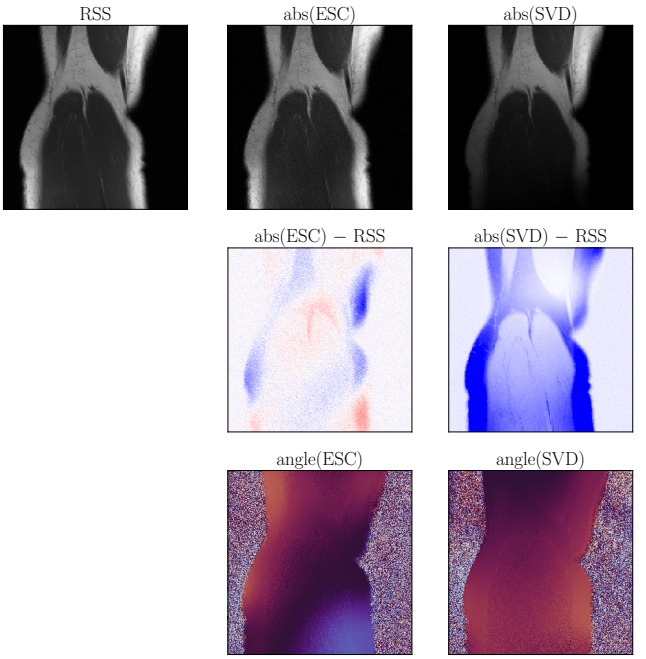

Figure 6. Fourth randomly chosen example.

Figures 3–8 show that the SVD-based compression to a single eigenmode for simulating the single coil yields dramatic differences from the ground-truth RSS, introducing severe artifacts. Compressing to only one eigenmode is apparently too drastic.

Figures 3–8 show that the reconstruction from ESC matches the ground-truth RSS reasonably well, but is noisy. The noise on the reconstruction from ESC is actually desirable for simulating a single-coil scanner (since a single-coil scanner is likely to be noisier than the ground truth estimated from the full multicoil data). This matches our expectation that ESC emulates a physically realizable process for converting a multicoil scanner into a single-coil scanner. The scanner which ESC simulates is genuinely physically realizable, even though it only approximates single-coil scanners that have actually been manufactured.

Figures 9–11 display examples for which ESC deviates substantially from the ground-truth RSS. These are among the worst problems we identified (though of course "worst" is somewhat subjective — SSIM of [12] was helpful in finding bad cases); about 4% of the volumes scanned exhibited some indication of such a problem, with about 10% of the cross-sectional slices within those volumes showing signs of the problems displayed in Figures 9–11. Overall, then, around 0.4% of the two-dimensional images in the data set display significant artifacts.

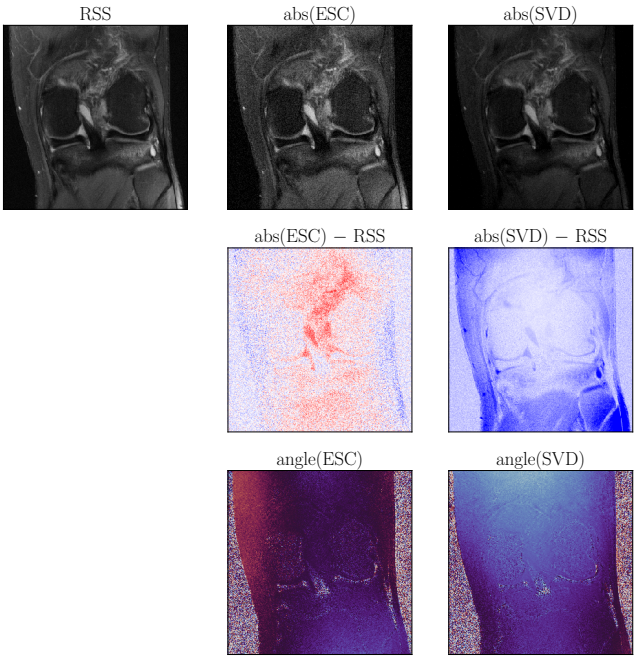

Figure 7. Fifth randomly chosen example.

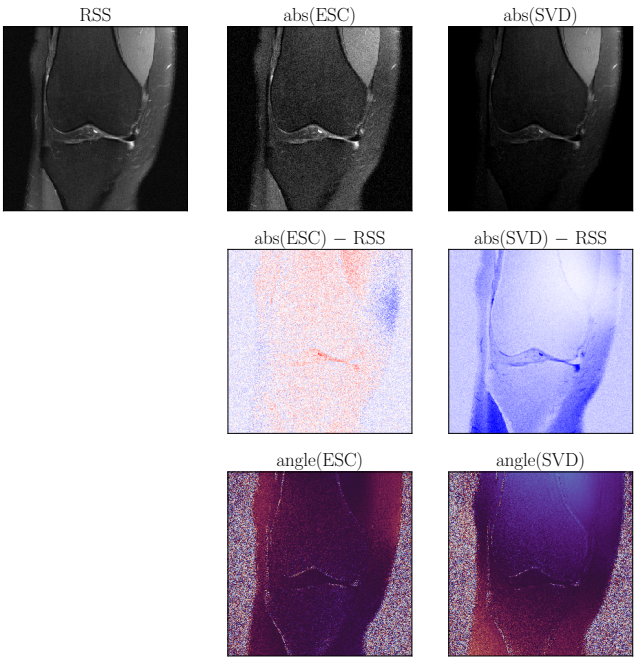

Figure 8. Sixth randomly chosen example.

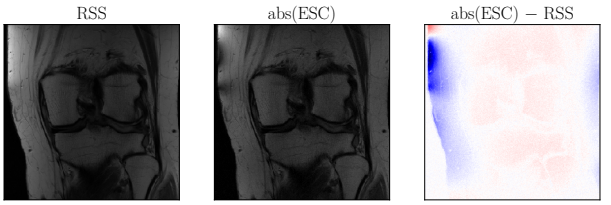

Figure 9. "Worst" abs(ESC)−RSS.

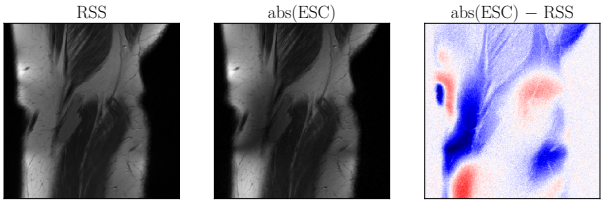

Figure 10. Second "worst" abs(ESC)−RSS.

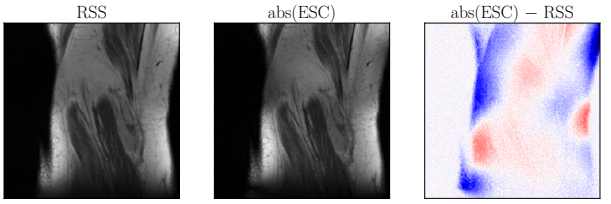

Figure 11. Third "worst" abs(ESC)−RSS.

The main culprit seems to be ringing artifacts near the boundary of the object being imaged. Whether such occasional ringing oscillations are an artifact of our fitting procedure or are inherent in the data when converting from multiple receivers remains unclear; however, the "ground-truth" RSS reconstructions do exhibit similar (albeit abated) artifacts. Having even a small number of cases with artifacts is indisputably undesirable, but then again data from the real world acquired in actual clinical practice is seldom ideal. Seeing as "ground-truth" RSS reconstructions exhibit similar (though less intense) artifacts, conducting simulations that avoid such artifacts may be impossible with this imperfect data.

Remark. The method proposed in the present paper is heavily data-driven. If computational simulations or excellent empirical estimates of the "physical coil sensitivity maps" defined in [5] could be calculated with high fidelity to the physical reality represented in the fastMRI data of [13], then less data-intensive methods might be feasible. Unfortunately, actual MRI scanners are exceptionally complex and deviate significantly from textbook idealizations; interactions of the objects

being imaged with the scanners further distort the sensitivities of receiver coils, as detailed, for example, by [4; 6]. While autocalibration procedures such as those from ESPIRiT of [9] work well for multicoil reconstruction, their estimates of sensitivity maps are specifically tailored for multicoil reconstruction, not for constructing single-coil measurements from multicoil.

4. Conclusion

Although our simulation of a single-coil scanner might appear perverse and pointless at first glance — after all, we throw away the information necessary to realize the gains in signal-to-noise ratios that imaging with multiple receivers provides — maintaining fidelity to a physically realizable simple MRI machine in fact yields a convenient test case, a reliable stepping stone toward the complexity of full parallel-MRI reconstruction. Less than 0.5% of the resulting emulated single-coil (ESC) images exhibit artifacts such as those displayed in Figures 9–11; Figures 3–8 display typical examples. The obvious alternative to our scheme would be to take measurements on an actually manufactured single-coil MRI scanner; however, MRI is extremely expensive, especially in the medical domain considered for the fastMRI data set of [13].

Acknowledgements

We would like to thank Matt Muckley, Erich Owens, Mike Rabbat, Dan Sodickson, and Anuroop Sriram.

Appendix: Entries of the solution vectors

Figures 12–14 display the entries of the solution x minimizing (1), where the ESC reconstruction of $|b|$ in (1) is $|Ax|$ from (1). No especially interesting patterns

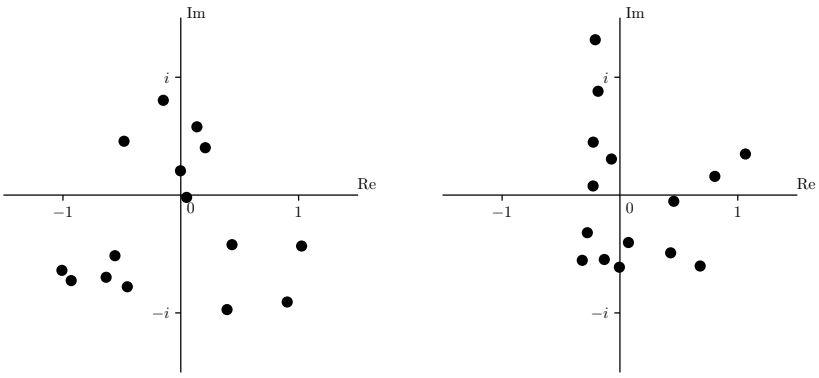

Figure 12. Entries of x minimizing (1) corresponding to the data in Figures 3 (left) and 4 (right).

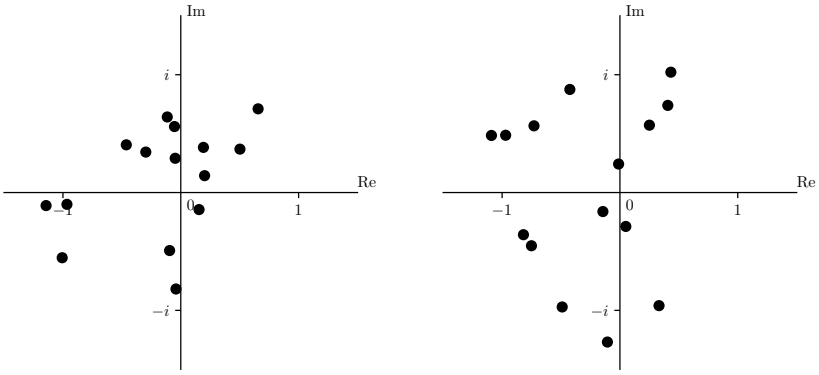

Figure 13. Entries of x minimizing (1) corresponding to the data in Figures 5 (left) and 6 (right).

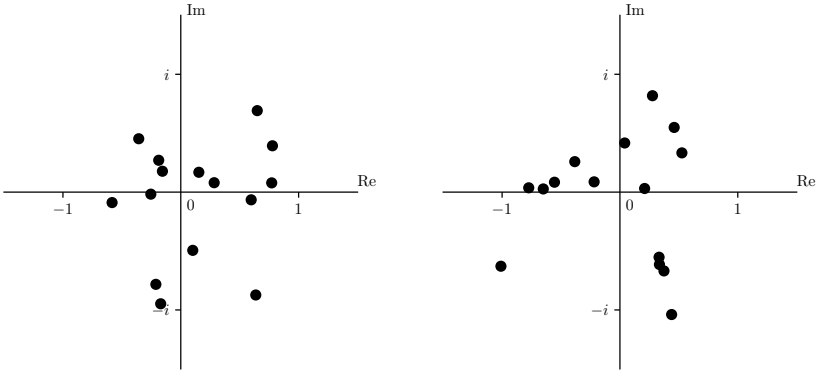

Figure 14. Entries of x minimizing (1) corresponding to the data in Figures 7 (left) and 8 (right).

appear to be evident; the coefficients seem to scatter somewhat randomly around the unit disk in the complex plane, not concentrating anywhere in particular.

References

[1] E. J. Candes and T. Tao, *Decoding by linear programming*, IEEE Trans. Inform. Theory **51** (2005), no. 12, 4203–4215. MR Zbl

[2] D. L. Donoho, *Compressed sensing*, IEEE Trans. Inform. Theory **52** (2006), no. 4, 1289–1306. MR Zbl

[3] E. Hellinger, *Die Orthogonalinvarianten quadratischer Formen von unendlichvielen Variabelen*, Ph.D. thesis, Georg-August-Universität zu Göttingen, 1907. Zbl

[4] J. H. Lee, J. M. Pauly, and D. G. Nishimura, *Partial k-space reconstruction for radial k-space trajectories in magnetic resonance imaging*, US patent US7277597B2, Stanford University, 2007.

[5] D. K. Ohliger, Michael A. Sodickson, *An introduction to coil array design for parallel MRI*, NMR Biomed. **19** (2006), no. 3, 300–315.

[6] J. M. Pauly, *Partial k-space reconstruction*, preprint, 2005.

[7] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller, *The NMR phased array*, Magn. Reson. Med. **16** (1990), no. 2, 192–225.

[8] K. Setsompop, L. L. Wald, V. Alagappan, B. A. Gagoski, and E. Adalsteinsson, *Magnitude least squares optimization for parallel radio frequency excitation design demonstrated at 7 Tesla with eight channels*, Magn. Reson. Med. **59** (2008), no. 4, 908–915.

[9] M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasanawala, and M. Lustig, *ESPIRiT — an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA*, Magn. Reson. Med. **71** (2014), no. 3, 990–1001.

[10] J. T. Vaughan and J. R. Griffiths (eds.), *RF coils for MRI*, Wiley, New York, 2012.

[11] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and contributors, *SciPy 1.0: fundamental algorithms for scientific computing in Python*, Nat. Methods **17** (2020), 261–272.

[12] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, *Image quality assessment: from error visibility to structural similarity*, IEEE Trans. Image Process. **13** (2004), no. 4, 600–612.

[13] J. Zbontar, F. Knoll, A. Sriram, T. Murrell, Z. Huang, M. J. Muckley, A. Defazio, R. Stern, P. Johnson, M. Bruno, M. Parente, K. J. Geras, J. Katsnelson, H. Chandarana, Z. Zhang, M. Drozdzal, A. Romero, M. Rabbat, P. Vincent, N. Yakubova, J. Pinkerton, D. Wang, E. Owens, C. L. Zitnick, M. P. Recht, D. K. Sodickson, and Y. W. Lui, *fastMRI: an open dataset and benchmarks For accelerated MRI*, preprint, 2019. arXiv

Received May 27, 2019. Revised May 1, 2020.

MARK TYGERT: mark@tygert.com
Artificial Intelligence Research, Facebook, Menlo Park, CA, United States

JURE ZBONTAR: jzb@fb.com
Artificial Intelligence Research, Facebook, New York, NY, United States

Communications in Applied Mathematics and Computational Science

msp.org/camcos

See inside back cover or msp.org/camcos for submission instructions.

The subscription price for 2020 is US $110/year for the electronic version, and $165/year (+$15, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues from the last three years and changes of subscriber address should be sent to MSP.

Communications in Applied Mathematics and Computational Science (ISSN 2157-5452 electronic, 1559-3940 printed) at Mathematical Sciences Publishers, 798 Evans Hall #3840, c/o University of California, Berkeley, CA 94720-3840, is published continuously online. Periodical rate postage paid at Berkeley, CA 94704, and additional mailing offices.

CAMCoS peer review and production are managed by EditFLOW® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

Communications in Applied Mathematics and Computational Science

vol. 15 no. 2 2020

1559-3940(2020)15:2;1-W