Learned Compression of High Dimensional Image Datasets

Elizabeth Cole[1], Qingxi Meng[1], John Pauly[1], Shreyas Vasanawala[2]
[1]Department of Electrical Engineering
[2]Department of Radiology
Stanford University
{*ekcole, qingxi, pauly, vasanawala*}@*stanford.edu*

Abstract

In many applications, such as burst photography and magnetic resonance imaging (MRI), multiple images are acquired to reduce the noise of the eventual reconstructed image. However, this leads to very high dimensional datasets which have redundant information across the various acquired images. In MRI, multiple images are acquired via multiple RF coil arrays in the scanner. Afterwards, coil compression is performed to convert the original set of coil images into a smaller set of virtual coil images to enable smaller datasets and faster computation time. However, traditional iterative coil compression methods are lossy and time-consuming. In this work, we propose a novel neural network-based coil compression method in pursuit of higher reconstruction accuracy and faster coil compression. Our learned compression method achieves up to 1.5x lower NRMSE and up to 10 times runtime speed compared to traditional methods on a benchmark test dataset.

1. Introduction

In many imaging applications, such as burst photography and magnetic resonance imaging (MRI), multiple images are acquired of approximately the same scene in order to reduce the noise of the eventual image. In MRI, this is done by using receiver arrays with multiple coil elements [17], enabling parallel imaging (PI) [9, 15, 18] acceleration. In parallel imaging, the known sensitivities and placement of the receiver coils are utilized to locate the signal, which allows the number of phase-encoding steps during signal acquisition to be reduced. Taking these multiple coil images and combining them into one image in the eventual image reconstruction process results in higher signal to noise ratio (SNR), where SNR increases as a factor of the square root of the number of images acquired. However, the large number of coils creates prohibitively large MRI datasets because these datasets contain multiple volumetric slices per patient and potentially a dynamic dimension. Additionally, a scan for a single patient typically contains multiple sequences (10-12 sequences per scan). The size of these datasets becomes increasingly problematic in terms of memory as well as infeasible computation time for reconstruction. Coil compression algorithms are effective in mitigating this problem by compressing data from many coils into fewer virtual coils.

Originally, coil compression was done in hardware [12]. However, such a hardware combiner does not consider the spatially varying sensitivities of the various coils, which leads to a larger loss in the compression of the entire signal [23]. Currently, there are two main traditional software methods for coil compression. The oldest and simplest method is to take the singular value decomposition [4] (SVD) across the set of various coil images. However, this method is slower compared to more advanced techniques. A faster method is geometric coil compression [23] (GCC). GCC removes the correlations between various coils across the fully-sampled readout dimension, allowing for reduction in data. Although GCC is faster than SVD, both are still iterative method, which are computationally expensive. In addition, the compression performance of both SVD and GCC drops off when the ratio of the dimensionality of original coils to virtual coils is very high.

Deep learning (DL) has been successful in a wide range of MRI applications, including reconstruction [5, 10, 13, 14, 16, 21] and segmentation [1]. DL has also been used for various types of data compression such as singular image compression [2, 7] and genomic data compression [6, 8]. In this work, we therefore propose and develop a neural network-based coil compression (NN-based CC) method to provide faster and more accurate compression on a *set* of images, rather than on the individual images themselves.

2. Methods

We propose a neural network-based framework to learn the coil compression task of mapping from our original coil space to our latent virtual coil space. Ultimately, our goal is to learn the *representation* of such a mapping. Inputs

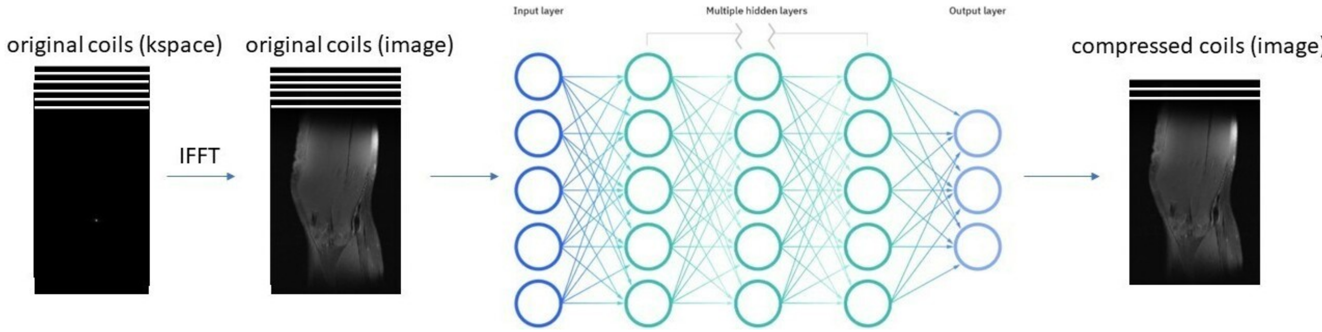

Figure 1. Framework overview of the rapid coil compression using neural networks. Our fully-connected neural network consists of 5 fully connected layers and it has 10 neurons per hidden layer. The number of channels in our input layer is equal to the number of original coils in our raw dataset. The number of channels in our output layer is equal to the number of virtual coils (the number of images in our set we wish to compress to).

to the network are the original coil images from the physical coil array. Outputs of the network are a set of virtual coil images. We can set the number of virtual coil images by setting the number of channels in the last layer of our network. We first apply the inverse fast Fourier transform (IFFT) to convert the original coil data from the frequency domain, otherwise known as k-space, to the image domain. Our raw data is in the frequency domain because the MRI signal is collected in the sparse frequency domain, not the signal domain. We separate the volumetric scan of the patient into slices along the axial dimension. Then, we feed multi-coil images, slice-by-slice, into a fully connected neural network. In the output layer of the neural network, we generate a set of virtual coils with a lower dimensionality than the original set of coils. This dimensionality is defined by the number of channels in the last layer of our network.

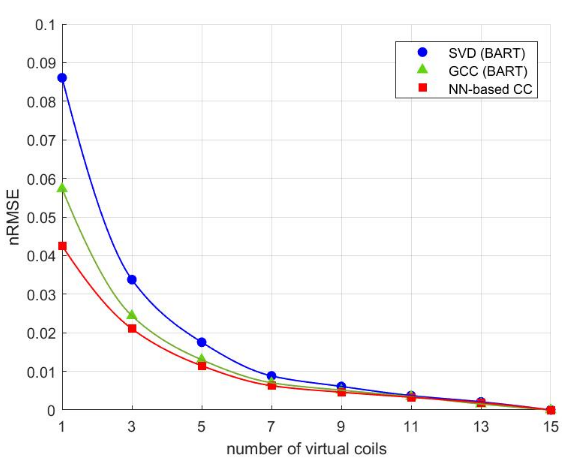

Figure 2. NRMSE vs. number of virtual coils for SVD, GCC and NN-based CC on the fastMRI test dataset.

We use the square root of sum-of-square (SSOS) to combine individual coil images into one image. The SSOS on an complex-valued image *m* with *n* coils is defined as:

$$SSOS(m) = \sqrt{\sum_{i=1}^{n}(|m_i|^2)} \tag{1}$$

where $|m_i|$ is the elementwise absolute value of each pixel in *m*.

We optimize the network parameters over the root mean square (RMS) loss between the SSOS image from the original coils and the SSOS image from the virtual coils. The

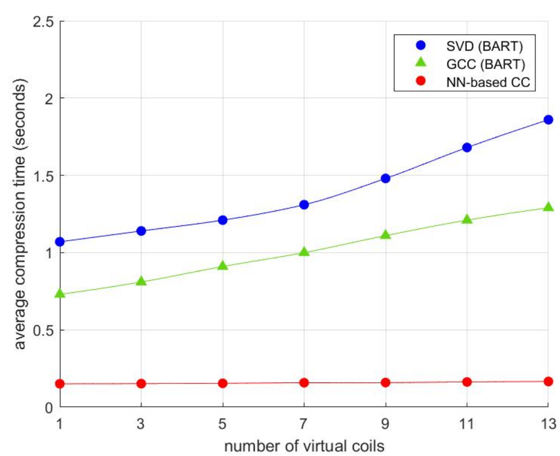

Figure 3. Average compression time vs. number of virtual coils for SVD, GCC and NN-based CC on the knee test dataset. NN-based CC was about 10 times faster than SVD and GCC. The compression time of our model was nearly the same for different numbers of virtual coils while the compression time for SVD and GCC increased as the number of virtual coils increased.

framework of our model is shown in Figure 1.

We train and test our model on the benchmark multi-coil knee dataset from fastMRI [22], which were acquired using a 15 channel knee coil array and a conventional Cartesian 2D TSE protocol on either 3T or 1.5 T scanners. Each subject volume contains slices of size 640x368 with 15 coils. There are 973 volumetric subjects with 34742 total slices during training and 56 volumetric subjects with 1959 slices during testing. Experiments were performed on a TITAN Xp GPU which had 64GB of RAM and 3TB of disk memory. We use normalized root-mean-square-error (NRMSE) between the original SSOS image and the compressed SSOS image to measure the compression accuracy. We compare our NN-based CC method with SVD and GCC, both implemented in BART [19], on the fastMRI dataset.

3. Results

Compression accuracy, measured as the NRMSE between SSOS images from original coils and virtual coils, is shown in Figure 2 for different numbers of virtual coils and different coil compression methods. Our NN-based CC consistently achieved smaller NRMSE compared to SVD and GCC. Figure 3 shows the compression time for SVD, GCC and our NN-based CC. Our NN-based CC was much faster than the traditional methods across all numbers of virtual coils. In Figure 4, we also visualize the SSOS image from virtual coils and its difference from the original SSOS image for SVD, GCC and our NN-based CC. Information contained in each virtual coil compressed by SVD, GCC and our NN-based CC is shown in Figure 5.

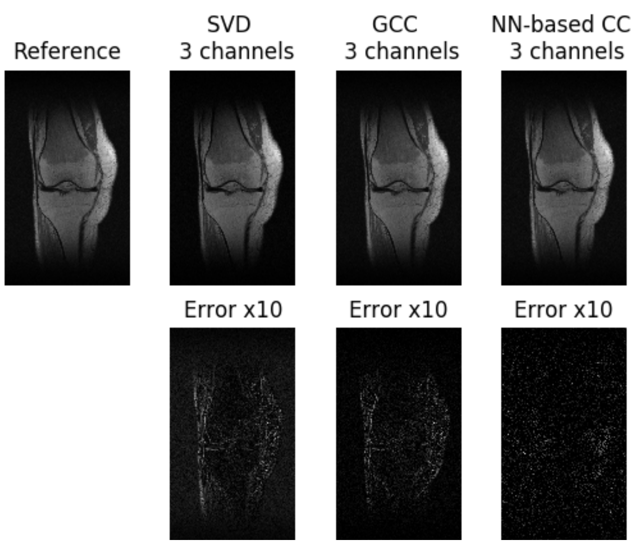

Figure 4. Comparison of SVD, GCC and NN-based CC. The original image (slice 1504 in the fastMRI knee training dataset) was used as the reference in the first column of the first row. The compression results of SVD, GCC and NN-based CC with 3 virtual coils were shown in the second, third and fourth column of the first row, respectively. Compression errors (10× in SVD, GCC and NN-based CC) were shown in the second row. The SSOS image from our model had smaller compression loss compared to both SVD and GCC.

4. Discussion

As shown in Figure 2, NN-based CC out-performs the traditional methods SVD and GCC. Also, it performs much better than the traditional methods when the dimensionality of the virtual coils is very small. Specifically, our method achieves more than 1.5x improvement in NRMSE over both SVD and GCC when the set of original coils is compressed to 1 virtual coil (i.e., a final composite image). As shown in Figure 3, another main advantage of the NN-based CC is that the compression time is much faster than the traditional methods. Although DL models take a long time to train - in our case, approximately 8 hours - training only needs to be done once. For our task of coil compression, inference time using our method is significantly reduced compared to the processing time of traditional iterative methods. One current large limitation is that a different model needs to be trained for every different dimensionality of virtual coils. For future work, we will explore generalizing one model to any dimensionality of virtual coils. Related to this model generalization is developing a model which is robust to dataset shift; i.e. does our model generalize well to the

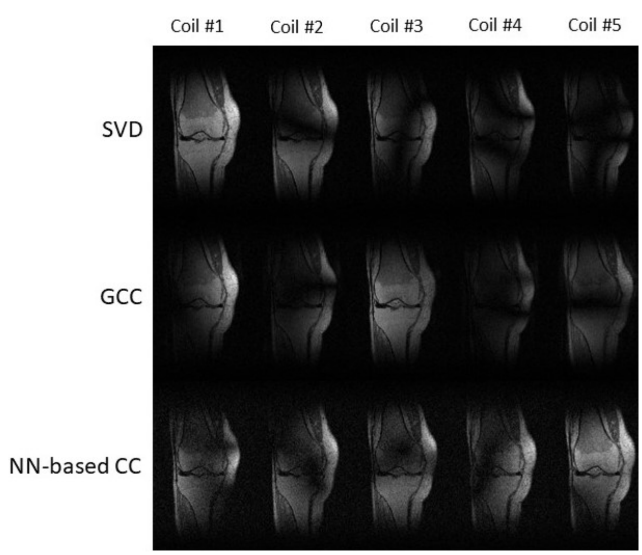

Figure 5. Magnitude of 5 virtual coils on a randomly chosen slice in the fastMRI knee training dataset for SVD (first row), GCC (second row) and NN-based CC (third row) respectively. For SVD and GCC, the information is concentrated in one or two virtual coils. For NN-based CC, the information was spread across each virtual coil, so it was more robust to the noise.

compression of image sets with anatomies/geometries that are different from those of the training data?

In this work, we only apply our method to Cartesian datasets; however, our method could also be applied to non-Cartesian datasets. There is actually a greater need for better coil compression in non-Cartesian datasets compared to Cartesian datasets because the current fastest iterative method, GCC, cannot be applied to non-Cartesian sequences [23]. This is because GCC attempts to reduce redundant information by splitting up the volume into separate subspaces based on slice-wise indexing. Additionally, some non-Cartesian sequences, such as 3D cones trajectories, produce even larger datasets compared to those produced by Cartesian sequences, due to the large number of coils used during the scan (approximately 60 coils). Our method could be also generalized to datasets with any number of original coils. For example, arrays with 32 coils are quite common for many types of anatomies, including knee, chest, and brain.

It would also be potentially useful and interesting to learn an *end-to-end* coil compression and image reconstruction framework, where both the compression and reconstruction are jointly learned. In this work, we use data which was fully-sampled in the frequency domain; however, ideally, we would want to subsample the data in the frequency domain to later leverage properties of parallel imaging during the reconstruction process. In such a learning framework, the ground truth of our model would be the SSOS image from fully-sampled data, while our inputs to the network would be subsampled coil images.

Additionally, using an encoder network to compress a set of images has implications in the computational photography domain under burst photography. In burst photography, the camera captures a set of consecutive images and they must be combined using some multi-frame super-resolution algorithm [3, 20]. Instead of storing all possible frames, which may be memory-expensive for mobile photography, multiple frames could be compressed into a virtual set of frames where the redundant information is removed, and only the unique information gathered by each frame remains. Such compression of multiple frames into a smaller set of virtual frames could be learned using our method. One important goal of most camera systems, especially in mobile photography, is immediate production of a photograph [11]. Our method fits well within this time constraint due to its improved speed over other traditional methods.

5. Conclusion

In this work, we propose a neural network-based framework to learn the representation of the coil compression task, where we wish to map from the space of a higher dimensional set of images to the space of a lower dimension set of images. To our knowledge, this is the first time MRI coil compression has been learned. Our learned representation achieves a higher compression accuracy compared to the two traditional state of the art iterative methods. Our method also achieves up to 10 times faster compression speed.

References

[1] Zeynettin Akkus, Alfiia Galimzianova, Assaf Hoogi, Daniel L. Rubin, and Bradley J. Erickson. Deep Learning for Brain MRI Segmentation: State of the Art and Future Directions, 2017. 1

[2] Pinar Akyazi and Touradj Ebrahimi. Learning-Based Image Compression using Convolutional Autoencoder and Wavelet Decomposition. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, volume 2019-June, 2019. 1

[3] Goutam Bhat, Martin Danelljan, Luc van Gool, and Radu Timofte. Deep Burst Super-Resolution. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2021. 4

[4] Martin Buehrer, Klaas P. Pruessmann, Peter Boesiger, and Sebastian Kozerke. Array compression for MRI with large coil arrays. *Magnetic Resonance in Medicine*, 57(6), 2007. 1

[5] Feiyu Chen, Valentina Taviani, Itzik Malkiel, Joseph Y. Cheng, Jonathan I. Tamir, Jamil Shaikh, Stephanie T. Chang, Christopher J. Hardy, John M. Pauly, and Shreyas S. Vasanawala. Variable-Density Single-Shot Fast Spin-Echo MRI with Deep Learning Reconstruction by Using Variational Networks. *Radiology*, 289(2):366–373, 11 2018. 1

[6] Wenwen Cui, Zhaoyang Yu, Zhuangzhuang Liu, Gang Wang, and Xiaoguang Liu. Compressing Genomic Sequences by Using Deep Learning. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 12396 LNCS, 2020. 1

[7] Thierry Dumas, Aline Roumy, and Christine Guillemot. Autoencoder Based Image Compression: Can the Learning be Quantization Independent? In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, volume 2018-April, 2018. 1

[8] Mohit Goyal, Kedar Tatwawadi, Shubham Chandak, and Idoia Ochoa. DeepZip: Lossless Data Compression Using Recurrent Neural Networks. In *Data Compression Conference Proceedings*, volume 2019-March, 2019. 1

[9] Mark A. Griswold, Peter M. Jakob, Robin M. Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 47(6):1202–1210, 2002. 1

[10] Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P. Recht, Daniel K. Sodickson, Thomas Pock, and Florian Knoll. Learning a variational network for reconstruction of accelerated MRI data. *Magnetic Resonance in Medicine*, 79(6):3055–3071, 6 2018. 1

[11] Samuel W. Hasinoff, Jonathan T. Barron, Dillon Sharlet, Ryan Geiss, Florian Kainz, Jiawen Chen, Andrew Adams, and Marc Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics*, 35(6), 2016. 4

[12] Scott B. King, Steve M. Varosi, and G. Randy Duensing. Optimum SNR data compression in hardware using an Eigencoil array. *Magnetic Resonance in Medicine*, 63(5), 2010. 1

[13] Alexander Selvikvåg Lundervold and Arvid Lundervold. An overview of deep learning in medical imaging focusing on MRI. *Zeitschrift fur Medizinische Physik*, 29(2):102–127, 2019. 1

[14] Morteza Mardani, Enhao Gong, Joseph Y. Cheng, Shreyas Vasanawala, Greg Zaharchuk, Marcus Alley, Neil Thakur, Song Han, William Dally, John M. Pauly, and Lei Xing. Deep Generative Adversarial Networks for Compressed Sensing Automates MRI, 5 2017. 1

[15] Klaas P. Pruessmann, Markus Weiger, Markus B. Scheidegger, and Peter Boesiger. SENSE: Sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 1999. 1

[16] Tran Minh Quan, Thanh Nguyen-Duc, and Won Ki Jeong. Compressed Sensing MRI Reconstruction Using a Generative Adversarial Network With a Cyclic Loss. *IEEE Transactions on Medical Imaging*, 37(6), 2018. 1

[17] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller. The NMR phased array. *Magnetic Resonance in Medicine*, 16(2), 1990. 1

[18] Daniel K. Sodickson and Warren J. Manning. Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine*, 38(4), 1997. 1

[19] Martin Uecker, Jonathan I Tamir, Frank Ong, and Michael Lustig. The BART Toolbox for Computational Magnetic Resonance Imaging. In *Proceedings of the 23rd Annual Meeting of ISMRM, Toronto, Canada*, 2016. 3

[20] Bartlomiej Wronski, Ignacio Garcia-Dorado, Manfred Ernst, Damien Kelly, Michael Krainin, Chia Kai Liang, Marc Levoy, and Peyman Milanfar. Handheld multi-frame super-resolution. *ACM Transactions on Graphics*, 38(4), 2019. 4

[21] Guang Yang, Simiao Yu, Hao Dong, Greg Slabaugh, Pier Luigi Dragotti, Xujiong Ye, Fangde Liu, Simon Arridge, Jennifer Keegan, Yike Guo, and David Firmin. DAGAN: Deep De-Aliasing Generative Adversarial Networks for Fast Compressed Sensing MRI Reconstruction. *IEEE Transactions on Medical Imaging*, 37(6):1310–1321, 6 2018. 1

[22] Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J. Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdzal, Adriana Romero, Michael Rabbat, Pascal Vincent, Nafissa Yakubova, James Pinkerton, Duo Wang, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastMRI: An Open Dataset and Benchmarks for Accelerated MRI. *arXiv*, pages 1–35, 2018. 3

[23] Tao Zhang, John M. Pauly, Shreyas S. Vasanawala, and Michael Lustig. Coil compression for accelerated imaging with Cartesian sampling. *Magnetic Resonance in Medicine*, 2013. 1, 4