Deep Reconstruction of Undersampled Cardiac MRI Datasets

Mario O. Malavé
Department of Electrical Engineering, Stanford University
350 Serra Mall, Stanford, CA 94305
momalave@stanford.edu

Abstract

Magnetic Resonance Imaging (MRI) acquisition and reconstruction can be a time consuming process which leads to delayed diagnosis by clinicians. One solution to this problem is by using undersampling which leads to more complex and extended iterative reconstruction times. To mitigate this problem, the reconstruction can be modeled using a convolutional neural network (CNN). In this work, we use a residual network (ResNet) to reconstruct undersampled cardiac datasets that were acquired using a gradient echo (GRE) sequence and a non-Cartesian trajectory. Qualitative evaluation on the test set suggests that undersampled reconstruction with the proposed model performs similarly compared to using the fully sampled k-space data.

1 Introduction

Scan time and reconstruction time is a key challenge for Magnetic Resonance Imaging (MRI). Scan time can cause discomfort for patients and long reconstruction times can lead to delayed diagnosis by clinicians. A typical MRI scan can last for several minutes depending on different parameter such as the resolution and field of view (FOV). Scan time can be decreased by using more time-efficient k-space sampling techniques and through the use of undersampling with compressed sensing [1]. This leads to an increase in the computation time required when reconstructing the undersampled k-space data. Deep learning has the potential for minimizing the reconstruction time of undersampled MRI data.

2 Related work

To solve the reconstruction problem when using undersampled MRI data, compressed sensing [1] is generally used to solve the problem. Recently, the previous iterative compressive sensing algorithm has been modeled using CNNs [2] [3]. In [2], the compressed sensing algorithm was reformulated as an unrolled optimization with deep priors (ODP) and in [3], sensitivity maps were used to reconstruct multi-coil MRI data. In this work, both approaches were used to reconstruct the undersampled multi-coil data.

3 Dataset

I acquired the cardiac dataset at the Magnetic Resonance Systems Research Laboratory (MRSRL) (which is part of the Department of Electrical Engineering at Stanford University) and at the Palo

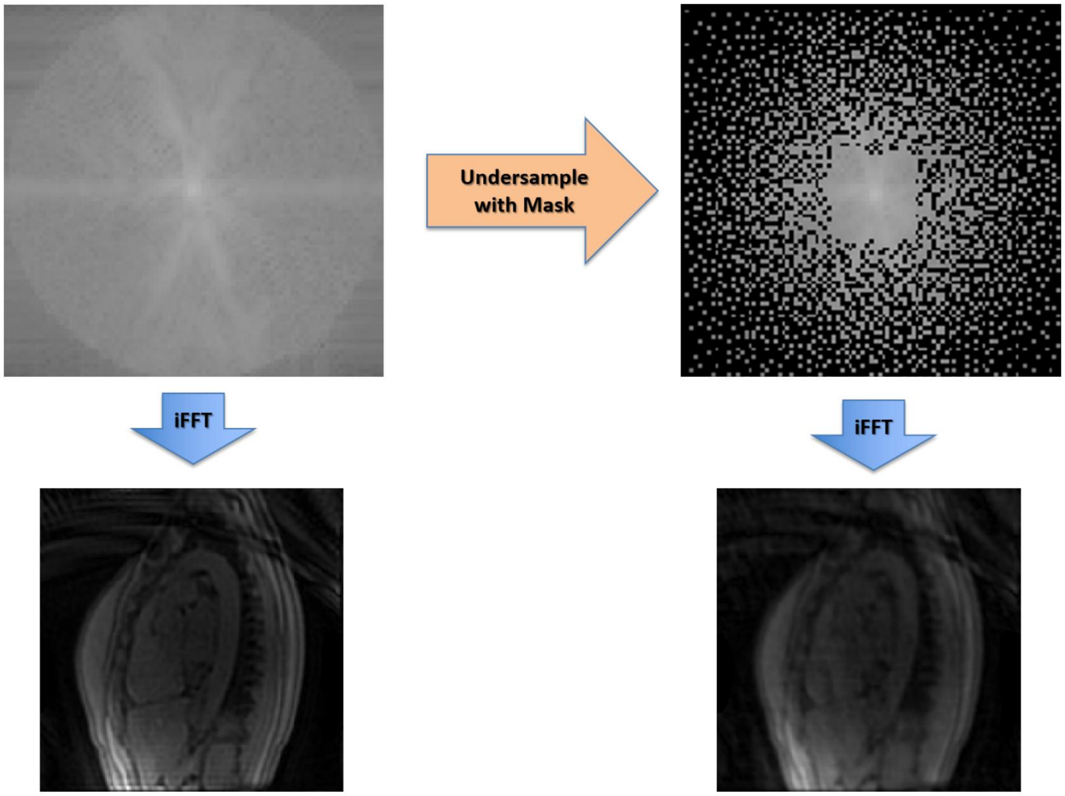

Figure 1: Example image before and after applying an undersampling mask.

Alto Medical Foundation (PAMF). The dataset comprised of 9,760 2D cardiac images (1/2 sagittal, 1/2 coronal slices) acquired using a gradient-echo (GRE) sequence using 12 spiral interleaves (75.6 ms) to achieve 28x28 cm^2 FOV and 3.1 mm in-plane resolution. We utilize 98% for training, 1% for validation, and the remaining 1% for testing.

The non-Cartesian 2D (navigator) data was first gridded to a cartesian grid. Then, the data was pseudo-randomly undersampled (by a factor of 1-2 in the x and y dimensions) using a variable density sampling mask which would normally be reconstructed using L_1-ESPRiT (compressed sensing) [1] [4]. This served as a method for data augmentation to expand the training set. In Figure 1, the k-space data, before and after one of the 9 different undersampling masks was applied, is shown with the corresponding iFFT images. The final step for data preparation included generating the coil sensitivity maps [4] for the 8 channels. This coil sensitives were used to combine the multi-coil data using SENSE reconstruction [5].

4 Methods

MRI data is acquired in the frequency domain (k-space) and is thus complex which means it has both real and imaginary components. There have been different approaches for solving this problem that include calculating complex weights, redefining the different activations functions, and CNN operations (i.e. batch normalization and max pooling). For this implementation, the real and the imaginary parts were separated into separate channels to handle the k-space data.

The current model architecture uses a residual network (ResNet) to solve the reconstruction problem similar to the "unrolled" design in [3]. The inputs into the ResNet are the undersampled k-space data and the respective 8 coil sensitivity maps for each channel. The input into each CNN is the coil combined image space data after performing SENSE reconstruction [5]. The CNN then uses 5 convolutional layers. Layers 1-4 use batch normalization, a rectified linear unit (ReLU) activation, and the final layer used a linear activation. Also, layers 2-5 use 128 features with a kernel size of 3x3, and the final layer is added to a skip connection from the input of the first convolutional layer to

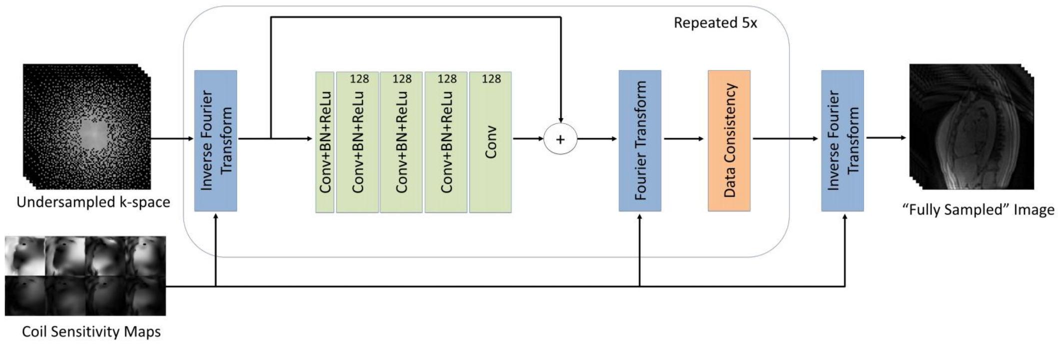

Figure 2: Implemented CNN architecture for image reconstruction.

accelerate training convergence. The data is then converted back to k-space (using the the Fourier transform and the coil sensitivity maps) for the data consistency step [4] and the process is repeated for five more iterations. A final inverse Fourier transform is then used to obtain the image. The architecture graph is shown in figure 2.

The network is trained using the complex l_1 loss (eq. 1) where x is the ground truth (fully sampled 2D cardiac images before applying the sampling mask) and $\hat{x}$ is the output of the network. When the fully sampled data is not available, the ground truth is the iteratively reconstructed image after using compressed sensing.

$$loss_{l_1} = \frac{1}{N} \sum_{i=1}^{N} ||x_i - \hat{x_i}||_1 \tag{1}$$

The architecture and objective function is trained using the Adam optimizer with a mini-batch size of 2. All the elements discussed are implemented in TensorFlow on a cluster with NVIDIA Tesla K80 graphics cards [6]. The overall model is trained for 10 epochs which took approximately 2 days.

5 Results & Discussion

When using a the complex l_1 loss, the training error (Figure 3) converged fairly quickly and significant improvements were seen in the reconstructed images for the validation and test sets. This was shown for both coronal and sagittal cardiac images. Output images are shown after testing the model in Figure 4 which display the input image (left), output images (middle) and ground truth (right). With the trained architecture, the cardiac images improved by recovering structures by essentially applying a denoising/smoothening operation. Performance was similar for larger undersampling ratios (2) compared to lower ratios (<2) in both x and y dimensions for the current architecture. Also, datasets that were collected for tracking the heart (navigators) gave similar motion estimates as the fully sampled counterpart.

6 Conclusion & Future Work

Qualitative and quantitative results of the ResNet "reconstruction" preserved structure and exhibited low complex l_1 loss. Other architectures may be warranted which decrease feature sizes for progressive convolutional layers and perhaps replace the iFFT/FFT blocks with fully connected layers. Also, the model hyperparameters can possibly be further tuned for improved performance. Future work includes expanding the results to 3D datasets (with 2D or 3D convolutions) and implementing the model using the full complex k-space data instead of separating it into multiple channels. Furthermore, expanding the training set to different MRI anatomies may be warranted to expand reconstruction performance of different scan locations.

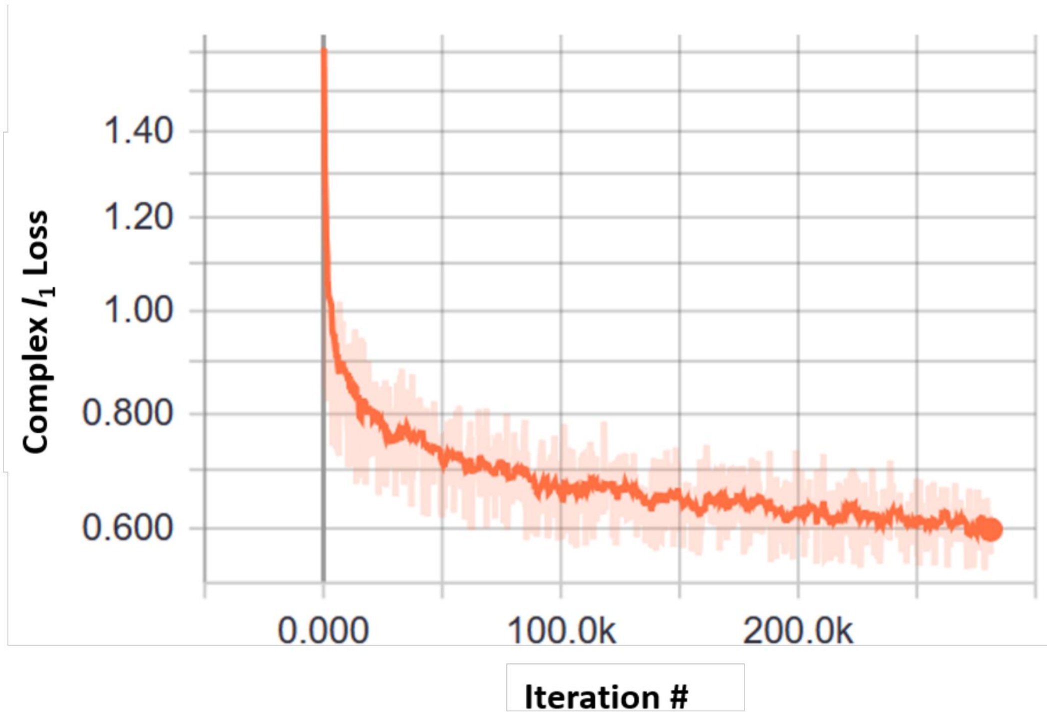

Figure 3: Complex l_1 error evaluation for each mini-batch when training the network.

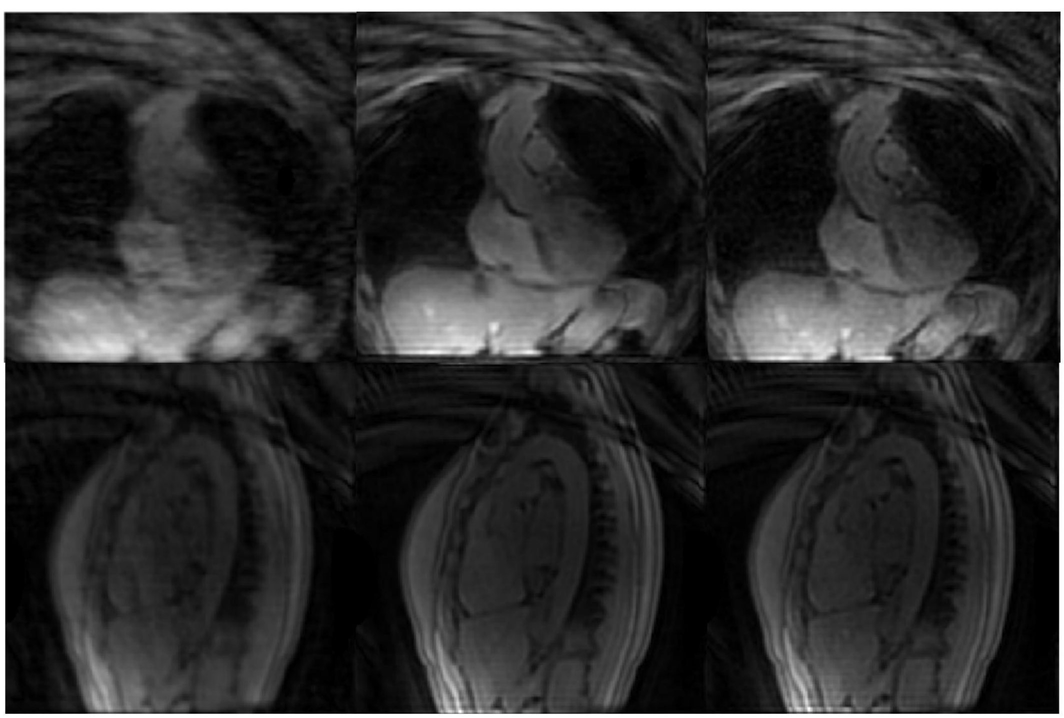

Figure 4: Example test coronal (top) and sagittal (bottom) datasets with the corresponding inputs (left), CNN outputs (middle), and ground truth (right) images.

7 Acknowledgements

This work is part of a research project being done in the Magnetic System Research Laboratory (MRSRL). I received help with the initial TensorFlow implementation from Chris Sandino and Joseph Cheng and had algorithm discussions with Srivathsan Koundinyan who are colleagues of mine in MRSRL.

8 Code

The project code is available at: `https://www.dropbox.com/s/i0cp8ci3armui0p/source_code%20-%20submit.zip?dl=0`

References

[1] Michael Lustig, David L Donoho, Juan M Santos, and John M Pauly. Compressed sensing mri. *IEEE signal processing magazine*, 25(2):72–82, 2008.

[2] Steven Diamond, Vincent Sitzmann, Felix Heide, and Gordon Wetzstein. Unrolled optimization with deep priors. *arXiv preprint arXiv:1705.08041*, 2017.

[3] 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*, 2017.

[4] Martin Uecker, Peng Lai, Mark J Murphy, Patrick Virtue, Michael Elad, John M Pauly, Shreyas S Vasanawala, and Michael Lustig. Espirit—an eigenvalue approach to autocalibrating parallel mri: where sense meets grappa. *Magnetic resonance in medicine*, 71(3):990–1001, 2014.

[5] Klaas P Pruessmann, Markus Weiger, Markus B Scheidegger, Peter Boesiger, et al. Sense: sensitivity encoding for fast mri. *Magnetic resonance in medicine*, 42(5):952–962, 1999.

[6] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283, 2016.