EE 451: Supervised Research Exposition

Anubhav Goel
Roll Number: 170040043
Department of Electrical Engineering
Indian Institute of Technology Bombay

Deep Learning Strategies for Reconstructing k-t Undersampled Resting State fMRI

Motivation

- Resting State fMRI measures low frequency fluctuations in BOLD signals
- These signals are used to determine the functional structure of the brain
- RS fMRI has quite a few medical applications such as presurgical planning for brain tumour and epilepsy patients and also possibly in the diagnosis of other neurological and psychiatric diseases
- The primary application of RS fMRI which is also the focus of this work is the determination of **Resting State Networks** which are discernable functional communities in the brain
- The primary need for undersampling of fMRI signals arises from the fact that undersampling can lead to better spatiotemporal resolution which in turn lead to improved brain connectivity maps
- Faster acquisition reduces the risk of data corruption due to patient motion and allows us to study dynamic networks

Review of Techniques

- The earliest technique which focussed on reducing the time for acquisition is called GRAPPA (Generalized Autocalibrating Partially Parallel Acquisition)
- This technique uses the idea of fully-sampled parallel imaging by making use of numerous magnetic coils and the properties of their spatial arrangement
- Other classical techniques rely on reconstruction using signal priors and the use of Bayesian graphical frameworks to learn data-adaptive prior models
- However, with the recent advances in deep learning and the advent of convolutional neural networks, there is a piqued interest in applying them to the problem of undersampled reconstruction
- After a brief review of classical techniques, we discuss the various neural networks that have been used so far and our own deep learning based approach to the problem

Dictionary Learning Approach

- This is one of the classical techniques that utilize the idea of learning a data-adaptive prior model through dictionaries that are designed to be robust to large physiological fluctuations that are typical in RS fMRI signals
- A Markov Random Field is used to model the 4D fMRI time series and the dictionary is represented as matrix of J columns where each column represents an atom of the dictionary
- The possible variations in the time-series underlying the functional maps is modeled as a similarity transform which consists of three parts:
  - A constant atom to model shifting
  - Coefficients to model scaling
  - An orthonormal transform to model rotation which is applied to all the non-constant atoms corresponding to one subject

Image Reconstruction Formulation

- In addition to the above, a heavier-tailed distribution than the Gaussian is used to model the physiological fluctuations and the quasi-norm of the dictionary coefficient vectors are penalized to enforce sparsity
- A maximum a posteriori estimate is learned over the training data of fully-sampled images to maximize the objective function below

$$P(x|D, A, R, \Phi) := \eta exp(-\sum_{v=1}^{V} \|x_v - \Phi_v(RDa_v + d_0 a_{v0}\|_2^p)$$

- The above objective is maximized by gradient ascent with adaptive step size that ensures the posterior probability increases at each step

K-t FASTER: Low Rank Constraints

- The k-t matrix obtained from the acquisition of MRI data can be modeled to low rank (barring the effects of thermal and physiological noise)
- The Matrix Completion problem is studied as an analogue to the Compressed Sensing problem where the data is represented in a known basis using a limited number of non-zero coefficients
- In Matrix Completion, the idea is to instead enforce low rank and not make use of a prespecified basis
- The problem is formulated in terms of a transform that takes the actual image to the undersampled k-space data
- A constrained optimization framework is used to obtain the reconstructed image as an optimal argument of the norm between the acquired data and the transformation of the argument under the constraint of a low-rank for the reconstruction

Hard Thresholding

- The above is achieved using the idea of hard thresholding that is iteratively applied to obtain an optimal reconstruction
- This drives all the singular values of the matrix (except a few) to zero, therefore enforcing the low rank prior on the result
- The reconstruction error is evaluated using Frobenius norm and validates the prior
- However, the iterative method is computationally inefficient and requires intensive computational power for the large k-t matrices which are common in MRI data

Application of Deep Learning

- With the success of the Deep Learning paradigm to various image processing and computer vision tasks, convolutional neural networks were also applied to the problem of undersampled MRI reconstruction
- One of the earliest architectures is simply a 3-layer convolutional network which learns the mapping between the undersampled image and the reconstruction by the use of training data in the k-space
- The architecture also corrects the k-space data which were present in the original acquisition but which were overwritten by the output of the neural network by taking a convex combination of the acquired value and the neural net output (a technique some of the later papers refer to as k-space correction)

$$\mathcal{F}u(k_x, k_y) = \begin{cases} S(k_x, k_y) & \text{if } (k_x, k_y) \notin \Omega \\ \frac{S(k_x,k_y)+\lambda S_0(k_x,k_y)}{1+\lambda} & \text{if } (k_x, k_y) \in \Omega \end{cases}$$

Early Architectures

- The architecture is quite simple consisting of vanilla convolutional layers and the weight initialization happens from a Gaussian distribution with zero mean and a standard deviation of 0.001
- The proposed network performs well and restores several details and fine structures to the reconstructed image
- The paper is one of the earliest works documenting the successful application of deep neural networks to the task of undersampled MRI reconstruction

Novel Architectures: UNET

- UNET, having successfully been used in several image processing tasks, was also applied to the task at hand
- The problem of undersampled reconstruction is again studied as an analogue to Compressed Sensing which is formulated as the following constrained optimization problem

$$y = argmin_y \|x - \mathcal{S} \circ \mathcal{F}(y)\|_{l_2}^2 + \lambda \|\mathcal{T}(y)\|_{l_1}$$

- The Deep Learning aims to learn a function from the undersampled image to the reconstruction through the training samples
- UNET provides a low dimensional latent space representation and the high dimensional features are recovered by concatenating feature maps from lower layers to the upper layers

Separability Problem

- Sub-Nyquist sampling is employed to speed up the time-consuming phase-encoding
- If we subsample by a rate of 2 in the phase-encoding, it causes a two-folded image when the inverse Fourier transform is taken (aliasing)
- The neural network needs to learn an unfolding map, however, an unfolding map will give the same result for two images which have an artifact that may only slightly differ in location from one image to the other
- This is called the separability problem

$$\mathcal{P} \circ \mathcal{S} \circ \mathcal{F}(y_1) = \mathcal{P} \circ \mathcal{S} \circ \mathcal{F}(y_2)$$

- The above problem can be addresses by adding some low frequency data which captures the numerous anomalies that form a crucial part of the MRI data and help in achieving separability in conditions such as above

Image Reconstruction

- The paper first investigates the zero-filled reconstruction (minimum norm solution) which is not practically useful in most cases
- The UNET is then used to filled the zeroes present at the unacquired locations to generate the reconstruction
- Finally, the acquired data may get distorted by the UNET, so it is overwritten by the original acquired data, a process which the paper refers to as k-space correction
- The mean MSE and mean SSIM are used as evaluation metrics which quantitatively validate the performance of UNET and also show the significant effects which are caused by k-space correction
- The high representativeness of the learned function is demonstrated by its flexibility as it gives good results on CT images as well on which it has never been trained

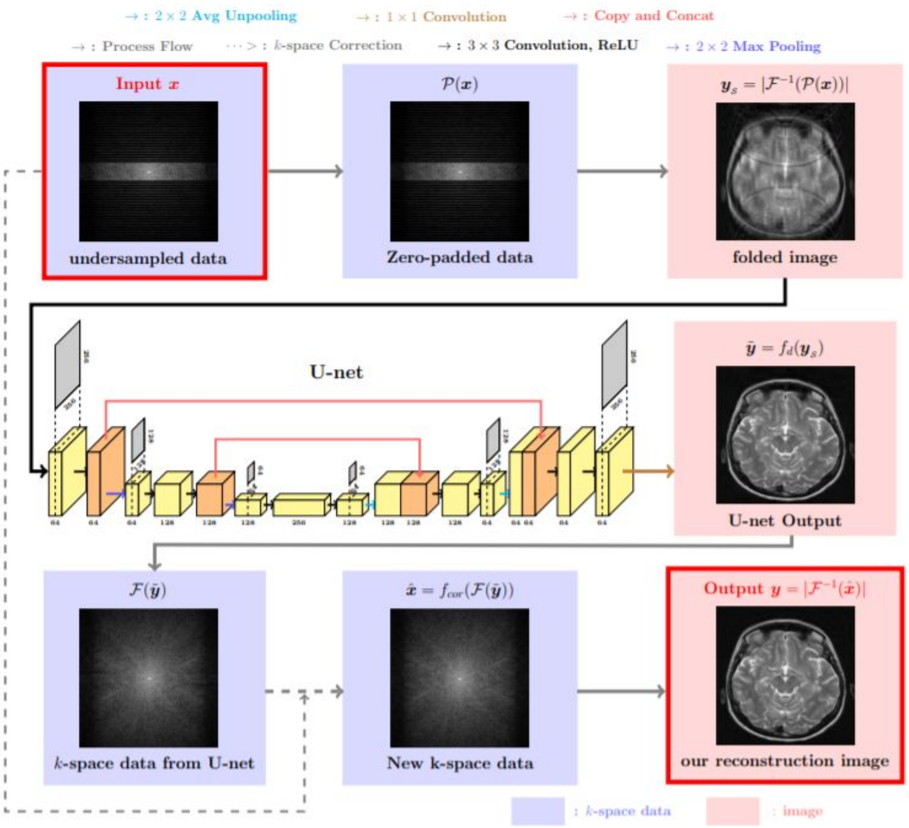

The architecture is summarized in this image. The zero padded k-space image is converted to spatial domain, passed through UNET and then k-space correction is performed. The UNET consists of two parts: the contracting network which performs down convolution and pooling, and the expanding part which performs upconvolution. The two parts are connected using skip connections.

Picture Credits: Chang Min Hyun, Hwa Pyung Kim, Sung Min Lee, Sungchul Lee, Jin Keun Seo, Deep learning for undersampled MRI reconstruction

Bayesian Approaches

- So far, the papers we have seen focus on the reconstruction of undersampled MRI images without any theoretical guarantees on the performance of these models
- The Bayesian approach models uncertainty associated with the performance of the model and the analysis is done over the following 2 kinds of uncertainty
  - Epistemic Uncertainty modeled by MC-dropout
  - Aleatoric Uncertainty modeled heteroscedastic loss
- Aleatoric Uncertainty refers to the irreducible uncertainty which is present in the data (measurement noise, aliasing pattern, artifacting)
- Epistemic Uncertainty refers to the fact that given a dataset there are more than one set of parameters for any given architecture which can faithfully reconstruct the data
- The above framework has been applied to the DC-CNN and UNET architectures

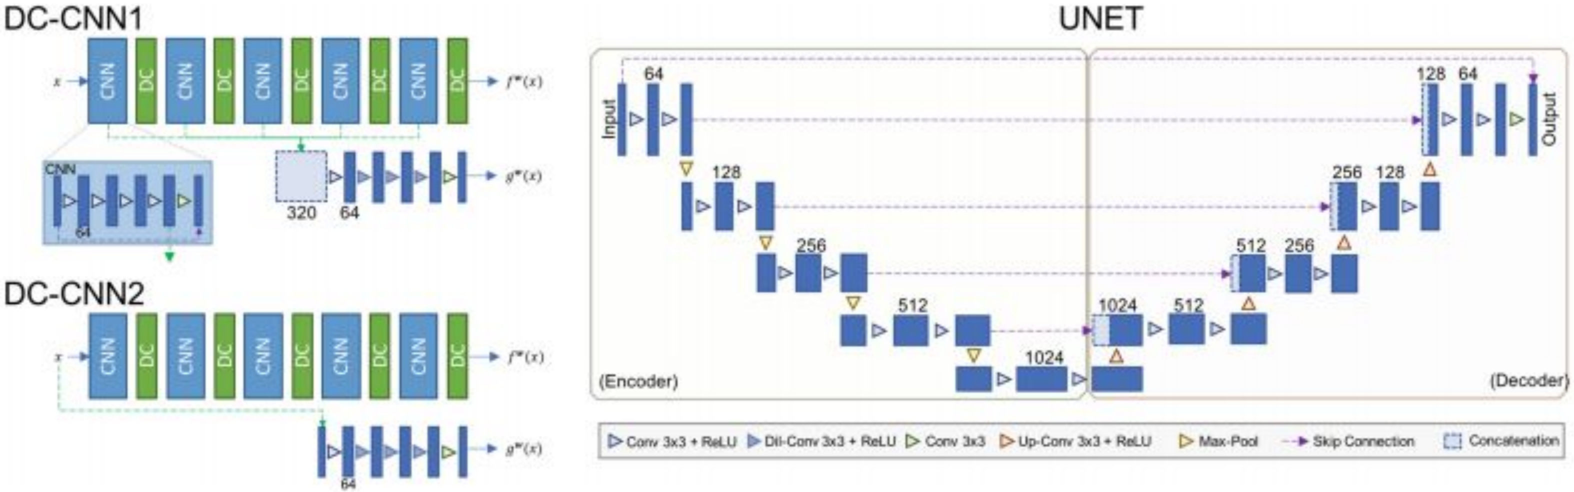

The architectures which were used by the Bayesian Deep Learning Paradigm. Notice the two variants of DC-CNN where the feature maps are taken from different locations to learn the uncertainty map.

Similarly, for UNET, the mean and the uncertainty map share the same encoder but a different decoder.

Picture Credits: Jo Schlemper, Daniel C. Castro, Wenjia Bai, Chen Qin, Ozan Oktay, Jinming Duan, Anthony N. Price, Jo Hajnal, Daniel Rueckert, Bayesian Deep Learning for Accelerated MR Image Reconstruction

Implementation and Analysis Work

- The work done after extensive literature survey on this topic focussed on further exploring the capabilities of Deep Learning frameworks in MRI reconstruction problem
- A three stage architecture is with end-to-end learning, which can be understood in the following 3 stages:
    - The first stage consisting of a CNN which fills in missing k-space data using the the data which has been acquired in the spatiotemporal neighbourhood
    - The second stage takes the data from k-space to spatial domain by the use of an inverse Fourier transform
    - The third stage uses a CNN learned for image quality enhancement in the spatiotemporal domain
- Uncertainty estimation makes the learning robust to large physiological fluctuations

Detailed Framework

- For the first stage of the network, the CNN receives the zero-filled k-space matrix for all time points for which acquisition has been done
- The images is divided into left and right halves and the set of weights for both these halves are trained separately
- The next stage takes the still temporally undersampled images which now have all the k-space entries filled and applies the inverse Fourier transform to them
- At this stage, the temporal interpolation is done (linear interpolation is done to keep the computational cost low) and this is passed on as input to the third stage
- A third stage of convolutional network which also involves residual mapping is used to model a Gaussian univariate wherein the mean represents the value of the pixel and the standard deviation gives an uncertainty map

Types of Losses

- Experiments were done with various types of losses which have been described as follows:
    - Mean-Squared Error Loss: This is the standard framework common in many deep learning paradigms, however, in our formulation we specifically also penalize the error with respect to the timepoints that were acquired before passing them to the third stage of neural network
    - Robust Loss: This takes into account the fact that RS-fMRI images can be corrupted with heavy physiological noise and a more heavy tailed distribution compared to the Gaussian is required. The mean-squared error loss is derived from the assumption of Gaussian distribution, so we penalize the p-th power of the Frobenius norm instead to model a heavier tail
    - Bayesian Loss: This maximizes the a posteriori probability of the reconstruction; given the mean and standard deviation obtained from Bayesian variant of the CNN

Results

- The results were measured using the mSSIM (mean structural similarity) and are summarized in the graph below
- The results were taken over 4 models which had been implemented: Bayesian Deep CNN, Robust Deep CNN, Deep CNN and Shallow CNN where the last two variants use the mean-squared error

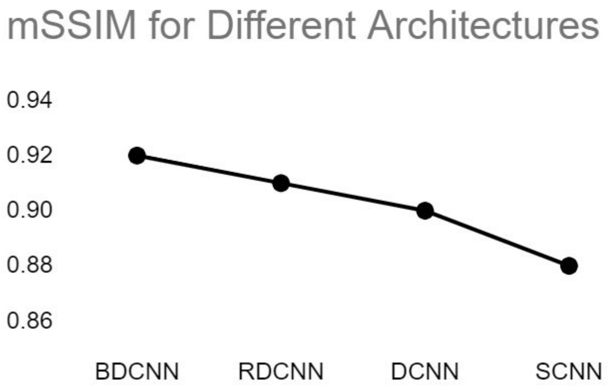

Future Work

- For the future, we are planning to explore the Expectation Maximization framework and integrate it into out setup
- Some preliminary reading regarding this has already been done
- Expectation Maximization is conventionally used in the unsupervised learning setup where in subsequent steps, the clustering centers are calculated and the a posteriori probability for a particular data point to belong to that cluster is maximized
- The posteriori probability is calculated on the basis of the estimate of a joint probability of latent space variables and the actual data
- This can be integrated into the reconstruction framework by working with the reconstruction problem in a latent manifold which turns out to be valid assumption in the case of RS-fMRI
- However, the setup is still raw and may require some further rigour before it can be used as formal problem-solving framework

References

- Mark A. Griswold, Peter M. Jakob, Robin M. Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, Axel Haase, Generalized autocalibrating partially parallel acquisitions (GRAPPA)
- Prachi H. Kulkarni, S. N. Merchant, Suyash P. Awate, Bayesian Reconstruction of R-fMRI from k-t Undersampled Data using a Robust, Subject-Invariant, Spatially-Regularized Dictionary Prior
- Mark Chiew, Stephen M. Smith, Peter J. Koopmans, Nadine N. Graedel, Thomas Blumensath, Karla L. Miller, k-t FASTER: Acceleration of Functional MRI Data Acquisition Using Low Rank Constraints
- Shanshan Wang, Zhenghang Su, Leslie Ying, Xi Peng, Shun Zhu, Feng Liang, Dagan Feng, Dong Liang, Accelerating Magnetic Resonance Imaging via Deep Learning
- Chang Min Hyun, Hwa Pyung Kim, Sung Min Lee, Sungchul Lee, Jin Keun Seo, Deep learning for undersampled MRI reconstruction
- Jo Schlemper, Daniel C. Castro, Wenjia Bai, Chen Qin, Ozan Oktay, Jinming Duan, Anthony N. Price, Jo Hajnal, Daniel Rueckert, Bayesian Deep Learning for Accelerated MR Image Reconstruction
- Karan Taneja, Prachi H. Kulkarni, S.N. Merchant, Suyash P. Awate, A Bayesian Deep CNN Framework for Reconstructing k-t-Undersampled Resting-fMRI
- All the equations have been taken from the report that accompanies this presentation