Magnetic Resonance Contrast Prediction Using Deep Learning

Cagan Alkan, John Cocjin, Andrew Weitz
{calkan, cocjinjb, aweitz}@stanford.edu

Background

Magnetic resonance (MR) images represent many different tissue contrasts depending on the acquisition paradigm that is used. Each contrast conveys specific information about the local tissue and its physical properties. Two of the most common MR contrasts are T1-relaxation and T2-relaxation.

Because the information provided by both T1- and T2-weighted images is important, clinical MR protocols will often collect both. However, MR acquisition is a slow and costly procedure. Predicting one MRI contrast from another would dramatically cut down the time and costs of clinical MR imaging.

In the last year, deep learning has been applied to several medical image transformation problems (e.g. MR to CT and MR to MR). These approaches have worked reasonably well, but also suffer from blurring effects in the output images.

Problem Statement

We propose the use of convolutional neural networks and deep learning to predict one MR contrast from another. We formulate the T1 to T2 contrast mapping as a regression problem and evaluate our results by comparing loss values.

Dataset

We are using the Human Connectome Project (HCP) data release of Mar 2017.

- T1 and T2 weighted 3D images from 1206 healthy subjects (3 Tesla scanner)
- Unprocessed data has size 320 x 320 x 256 per contrast per subject
- Processed data has size 260 x 260 x 311 per contrast per subject
- Six binary tissue masks for each image (gray matter, white matter, etc.)

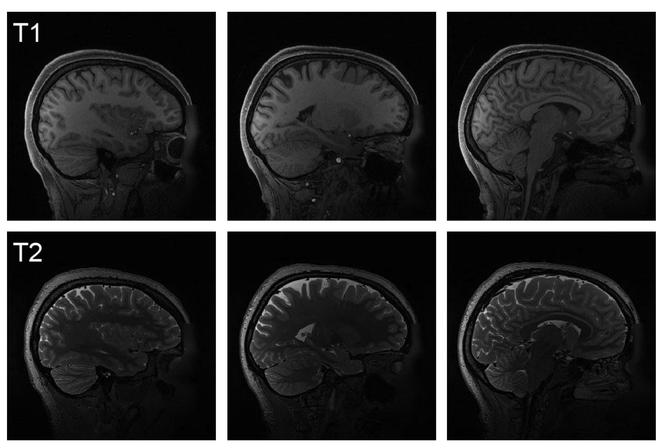

Models

We implement an eight-layer convolutional neural network. As input, we use 2D patches of size 128x128 of the T1-weighted images concatenated with the corresponding masks of six different tissues types. Thus, each input is a 3D matrix of size 128x128x7. The first seven layers of the network consist of a 2D convolution operation, spatial batch normalization, and a ReLU activation function. Each 2D convolution uses 32 filters of size (3,3). The final layer includes a 2D convolution with 1 filter to collapse the image into the desired output size. Training and model fitting were implemented in Python using Keras with a TensorFlow backend. Gradient descent was performed with the Adam method using default hyperparameters. All models were trained using a 80-2-18 breakdown of the data into training, validation, and test sets. Twenty epochs were run.

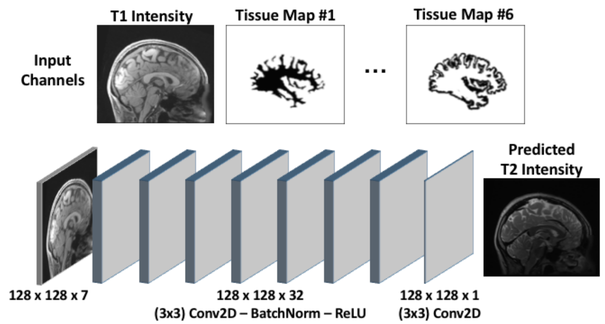

Experimental Evaluation and Findings

To compare output T2-weighted images with ground truth T2-weighted images, we use the mean square error function (i.e. L2 difference), shown in Equation 1.

$$d_2(I_1, I_2) = \frac{1}{N}\sum_{n=1}^{N}(I_1[n] - I_2[n])^2 \quad (1)$$

We consider two different inputs:

- 2D T1-weighted input images with no tissue information
- 2D T1-weighted input images with six additional tissue mask channels

The final loss values after twenty epochs for each method are presented in Table 1. We find that adding tissue mask labels to the input lowered loss values in both training and test sets by approximately 10%. Furthermore, we find that using the model trained with one subject's data on another subject's test input gives only slightly higher loss values.

Table 1: Loss value comparisons.

Model	Dataset	Loss
No Tissue Masks	Training	8,529
No Tissue Masks	Within-subject Test	8,548
No Tissue Masks	Between-subject Test	13,175
With Tissue Masks	Training	7,606
With Tissue Masks	Within-subject Test	7,602
With Tissue Masks	Between-subject Test	11,610

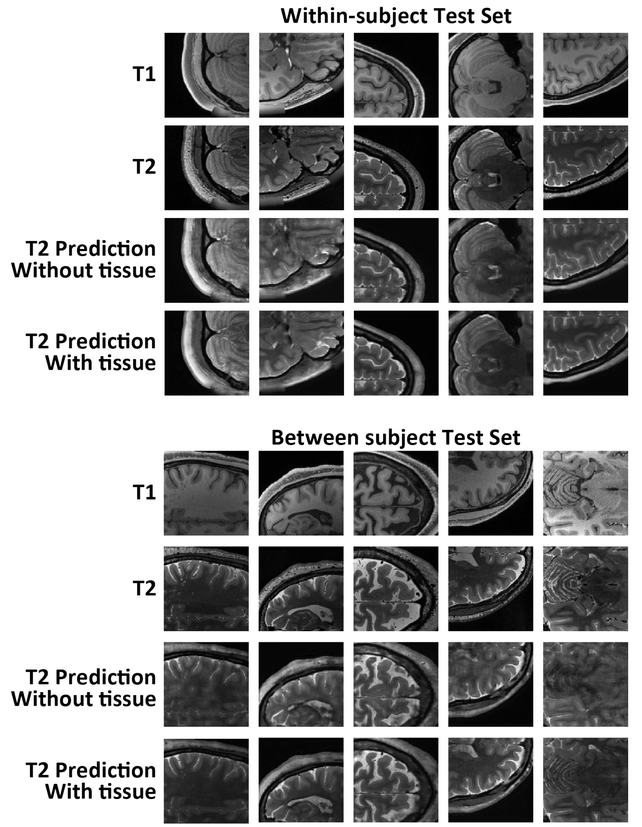

Conclusion and Future Directions

Our results demonstrate the feasibility of using deep convolutional neural networks to predict one MRI contrast from another. We will also: (1) Use 3D convolutions to better utilize local spatial information, (2) Incorporate residual blocks, (3) Evaluate model's ability to properly map T1 to T2-weighted images at site of abnormalities.