Deep Residual Nets for Improved Alzheimer's Diagnosis

Aly Valliani
Swarthmore College
aavalliani@gmail.com

Ameet Soni
Swarthmore College
soni@cs.swarthmore.edu

ABSTRACT

The field of image analysis has seen large gains in recent years due to advances in deep convolutional neural networks (CNNs). Work in biomedical imaging domains, however, has seen more limited success primarily due to limited training data, which is often expensive to collect. We propose a framework that leverages deep CNNs *pretrained* on large, non-biomedical image data sets. Our hypothesis, which we affirm empirically, is that these pretrained networks learn cross-domain features that improve low-level interpretation of images. We evaluate our model on brain imaging data to show our approach improves the ability to diagnose Alzheimer's Disease from patient brain MRIs. Importantly, our results show that pretraining *and* the use of deep residual networks are crucial to seeing large improvements in diagnosis accuracy.

ACM Reference format:
Aly Valliani and Ameet Soni. 2017. Deep Residual Nets for Improved Alzheimer's Diagnosis. In *Proceedings of ACM-BCB '17, Boston, MA, USA, August 20-23, 2017,* 2 pages.
https://doi.org/10.1145/3107411.3108224

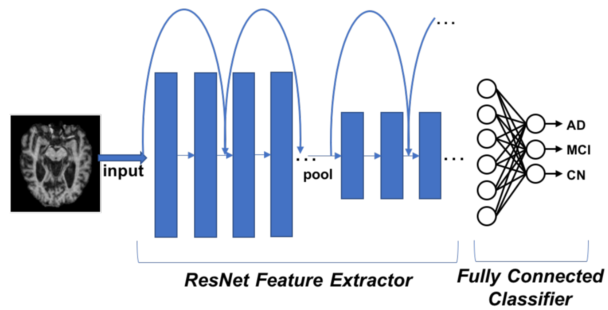

Figure 1: An overview of our approach. ResNet Feature Extractor refers to the pretrained 18-layer residual network defined by He et al. [2]. Each layer defines a typical 3x3 convolution layer (not all layers are shown) with the addition of residuals (arcing arrows across layers). The last two layers are a fully connected neural network trained from scratch for AD classification.

1 INTRODUCTION

Alzheimer's Disease (AD) is a neurodegenerative disorder affecting over 5.3 million people in the United States. Diagnosis by experts is difficult, usually occurring much after symptoms have set in, and can only be verified via a postmortem autopsy. Compounding matters is that early signs of Alzheimer's are difficult to differentiate from mild cognitive impairments (MCIs), which are often a result of aging rather than onset of disease. While promising, brain imaging remains an underutilized resource for aiding medical experts in performing early diagnosis due to limitations of the human eye. Automating this analysis for decision support is itself hindered by the limited size of image data sets to help train models, particularly deep convolutional neural networks(CNNs) [3] which have shown promise in many imaging problems. This limitation exists in many bioimaging tasks, where the data is too expensive or sensitive to generate in large quantities.

Suk and Shen [5] first proposed deep learning approaches for AD diagnosis, utilizing a sparse autoencoder with a multi-modal SVM to combine MRI and PET images from a patient. Gupta et al. [1] showed improvements through the use of a single-layer CNN pretrained on a small set of natural images. Payan and Montana [4] further extended this framework to 3D CNNs. The main shortcoming of existing CNN-based approaches is that the networks are shallow, with only a single layer of convolutions to learn a latent feature representation of the data. This limitation prevents the learning of hierarchical representations of the data, which is crucial to medical tasks where morphological changes are often subtle and multifaceted. Learning deep networks, however, is difficult due to vanishing gradients and the need for very large training sets.

We propose the use of *pretrained residual network* models for predicting Alzheimer's Disease from brain images. Specifically, we utilize the the ResNet [2] network which finished atop the 2015 ILSVRC ImageNet competition. This network is trained on millions of natural images, thus overcoming the limitation of data size. Second, the residual architecture allows for the learning of "very deep" networks, which are empirically more accurate and easier to optimize.

We validate our hypothesis that pretrained deep residual networks improve AD diagnosis by performing 3-way classification (AD vs. MCI vs. healthy) on brain MRIs provided by the Alzheimer's Disease Neuroimaging Inititative (ADNI). Our results show that deeper and pretrained neural networks surpass shallower networks in classification accuracy.

2 APPROACH

Convolutional neural networks are hierarchical methods for end-to-end feature learning. They are composed of a series of nonlinear functions that transform pixels from an input image into class scores for prediction. Convolutional layers break the input into *receptive fields* where weights are tied across fields – preventing

ACM-BCB '17, August 20-23, 2017, Boston, MA, USA

ACM ISBN 978-1-4503-4722-8/17/08.
https://doi.org/10.1145/3107411.3108224

overparametization as in multilayer perceptrons. The layer-wise architecture enables the network to learn increasingly abstract spatial features to differentiate object categories.

The residual neural network is a CNN variant that employs shortcut connections (as seen in Fig. 1) to allow input from lower layers of the network to be available to nodes at higher layers. These connections are constructed from residuals blocks that approximate a residual function using input transformed from the preceding layer and identity mappings of input from layers much further down the network. Specifically, if a typical layer aims to learn a latent representation $F(x)$, a residual block models this representation with the inclusion of an identity connection; i.e., $H(x) = F(x) + x$ (see He et al. [2] for details). The architecture allows multiple pathways for gradients to flow through the network, which permits the creation of much deeper networks without the burden of vanishing gradients. The residual blocks have also been more recently conceptualized as independent networks, thereby making the residual network an ensemble of multiple independent networks.

For our task, we construct a deep residual network consisting of 18 layers modeled after the ResNet-18 architecture. To initialize weights, we use the already learned weights on the ImageNet data set as specified in He et al. [2] (thus, the term *pretrained*). Since our task is different than ImageNet, we only take the convolutional layers as filters (i.e., we omit the fully connected classifier). We add two fully-connected layers, with 1000 and 100 hidden units respectively, that predict three outputs using a softmax classifier. The pretrained network is fine-tuned on MRI (see below) with real-time data augmentation (affine transformations) to prevent overfitting. All networks include batch normalization after every convolutional layer and utilize the ReLU activation function. Networks were trained with mini-batch stochastic gradient descent using an early-stopping criteria.

3 RESULTS AND DISCUSSION

Data used in the preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.ucla.edu). Our data set includes the median axial slice of 660 ADNI images (only the first image for each patient was maintained to avoid data leakage), 188 of which were diagnosed as having Alzheimer's Disease (AD), 243 as Mild Cognitive Impairment (MCI) and 229 as Cognitively Normal (CN) based on examination by a medical expert. Each image was skull-stripped and registered. Using an 80/20 train/test set split, we aimed to assess our hypothesis that pretrained residual networks would improve AD diagnosis. To do so, we asked the following questions:

- **Q1**: Does a pretrained residual network transfer to the MRI domain to improve prediction in AD diagnosis?
- **Q2**: Does pretraining influence ResNet's success?
- **Q3**: Does data augmentation improve the ResNet's ability to adapt to MRI images?

To answer these questions, we compare four different classifiers using accuracy on predicting on the 2-class AD vs. CN problem as well as the more difficult 3-way classification (AD vs. MCI vs. CN). The results on the held aside test set are in Table 1. All approaches are implemented using the Torch7 library.

Model	AD vs. CN	3-way
Baseline CNN	73.8%	49.2%
ResNet	77.5%	50.8%
Pretrained ResNet	78.8%	56.1%
Pretrained ResNet + augmentation	**81.3%**	**56.8%**

Table 1: Accuracy on Alzheimer's Disease (AD) vs. Cognitively Normal (CN) classification and 3-way classification (AD vs. MCI vs. CN)

For **Q1**, we trained two networks – the proposed approach in Sec. 2 (pretrained ResNet + augmentation) and a baseline CNN of one convolutional layer containing 5x5 kernels and 64 feature maps, and two fully-connected layers containing 1000 and 100 hidden units, respectively, each with dropout prior to the non-linearity (to approximate existing approaches [1, 4]). Our results show a large improvement in accuracy over the baseline CNN model on both tasks; thus we can answer **Q1 affirmatively** – the ResNet structure successfully adapts to the MRI domain and improves prediction.

Q2 asks whether this improvement is due to pretraining, the deep residual structure, or both? Our results show that both features are important. The ResNet using randomly initialized weights improves upon the baseline CNN in both tasks. Furthermore, pretraining boasts higher accuracy on both tasks than the randomly initialized ResNet. Thus, we can answer **Q2 affirmatively** that both are key aspects to the result, though with different magnitudes of importance in the two tasks we tested.

Lastly, a key method for regularizing networks and simulating more data is the use of real-time data augmentation (affine transformations of the data through rotations, flips and translations during training). The last two rows in Table 1 examine a pretrained ResNet without and with data augmentation, respectively. We can also answer **Q3 affirmatively**, as augmentation improves accuracy on both tasks.

The results of this initial work show that our framework makes significant contributions through the use of both pretraining and very deep residual neural networks. As part of future work, we hope to understand more closely the contribution of pretraining versus depth. Additionally, we are currently extending the network to use 3D convolutions and exploring other avenues for volumetric analysis. Finally, we plan to transfer our model to the more important medical question of early diagnosis: can we predict which patients with MCI are likely to later develop Alzheimer's Disease.

REFERENCES

[1] Ashish Gupta, Murat Ayhan, and Anthony Maida. 2013. Natural Image Bases to Represent Neuroimaging Data. In *Proceedings of the 30th International Conference on Machine Learning*. Atlanta, Georgia, USA, 987–994.
[2] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep Residual Learning for Image Recognition. *CoRR* abs/1512.03385 (2015).
[3] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
[4] Adrien Payan and Giovanni Montana. 2015. Predicting Alzheimer's disease: a neuroimaging study with 3D convolutional neural networks. *CoRR* abs/1502.02506 (2015).
[5] Heung-Il Suk and Dinggang Shen. 2013. Deep Learning-Based Feature Representation for AD/MCI Classification. In *Proceedings of the 16th International Conference on Medical Image Computing and Computer-Assisted Intervention*. 583–590.