

Project Proposal

Investigating effectiveness of various CNN architectures for medical imaging with a focus on Alzheimer's studies

Subject Specialist: Gordon Dickers

This paper lays out a proposal for an experiment looking at the potential effectiveness of convolutional neural network (henceforth referred to as CNN) architectures to identify Alzheimer's based on MRI datasets. To look at this it first outlines a basic history of different CNN architectures as they have become more advanced and then goes on to look at what potential options there use for use in the experiment. The specific case for the experiment is then discussed, going in to detail about what the plan for conducting the experiment is and what exactly will be considered a successful condition for the experiment; this includes where the data to be used for the training will be gathered from, what constants shall aim to be maintained, and what the variable factors in the experiment could be.

Table of Contents

Introduction	3
1 – History of deep neural networks	3
2 – Previously proposed use cases of CNNs for medical imaging.....	4
3 – Proposed experiment based on MRI datasets.....	5
4 – Research and development methods	6
5 – Resources and planning.....	6
References	8

Introduction

The main aim of the project this proposal is discussing is to experiment with CNN architectures and build a stronger understanding of how such networks learn, to aid in further study at postgraduate level. The area of interest for the experiment is the human brain, so the intention at this stage is to use open source medical data, namely MRI scans from Alzheimer's studies, to train networks. This proposal therefore aims to look at previous studies that have been carried out regarding the use of medical data to train CNNs; this should include what specific architecture was being used, what type of data was included in the network training datasets, and how successful the experiment was considered to be. It will discuss which sensible objectives have been identified to constitute a successful project and a more detailed plan of how the experiment shall be carried out. It is also of importance to discuss any ethical issues that may arise from the experiment being carried out, in this case the use of medical data.

1 – History of deep neural networks

To carry out an experiment making use of CNN architectures, it is important first to understand the progression that has been made in the structure and complexity of available architectures. All but one of the architectures to be covered have been part of the ImageNet competition [1] (the oldest network not being part of it as the competition did not exist at that point).

In this brief research, the oldest example to be looked at is LeNet-5, which aimed to use back-propagation algorithms as a form of gradient-based learning in order to identify writing numbers; mainly as part of banking systems for recognising amounts of money on cheques [2]. This network architecture took 32x32 pixel grayscale images and consisted of 7 layers. This network itself technically is not an example of a deep network, as the 7 layers meant it was very shallow (in order for the architecture to work with the technology available at the time); resulting in the network only being able to take low resolution grayscale images as input [3].

Next is AlexNet [4] from 2012. AlexNet was trained to classify images between 1000 different classes using a dataset of over 1.2 million images. The architecture won first place in the 2012 competition [1], achieving a better accuracy rate than any other architecture for the given dataset. This architecture managed to achieve this high accuracy by making use of the (at the time) “recently developed regularization method called ‘dropout’”. Dropout is a rather interesting method of reducing the likelihood that an architecture will end up overfitting to the dataset (overfitting meaning that the network will have a high accuracy on the training and test datasets, but a much lower accuracy on other images – essentially the network learns too specifically about the datasets images and not the patterns it should be looking for). A way of reducing overfitting is to have collections of networks with different configurations working together, this is expensive however as it requires more time and computing power to train, and means multiple models need to be maintained over time. What dropout does is simulate having multiple models by randomly dropping nodes during the training process [5] [3]. Certain layers can be chosen to be dropout layers in architectures.

ZFNet came in 2013, though it does not have much that is different to note as it is a very similar architecture to AlexNet. The changes made, however, did mean that ZFNet managed to get a higher accuracy rate than AlexNet. The main aim of their study with this architecture was to further understand how AlexNet managed to get such high accuracy and look at what is happening in terms of calculations at individual layers [6] [3].

Next is GoogLeNet, which was originally codenamed 'Inception', from 2014. It managed to achieve a much higher accuracy than AlexNet, while at the same time reducing the number of parameters in the architecture from 60 million to 4 million, significantly reducing the training time required [7] [3]. The main goal of this project was to keep the computational cost of the network as low as realistically possible, reducing as much computational waste as possible. This was primarily achieved by making efficient use of densely packed sparse-matrix computations. The name 'Inception' is still being used as the name of the Google architecture, with GoogLeNet being the name given to the specific entry that was used in the competition. There have been multiple versions of Inception since, with the most recent being Inception v4.

The last one to look at is ResNet. The goal of the ResNet architecture is to make optimization easier using residual learning techniques. The development team noted that in examples of very deep architectures a point would be reached where increased depth causes the accuracy to either stagnate or even degrade. Through the use of residual learning, the specific functions at each layer are defined more explicitly and are therefore easier to optimize, allowing more residual layers to be stacked for increased depth without degradation, and potentially without stagnation [8] [3].

2 – Previously proposed use cases of CNNs for medical imaging

As CNNs get progressively more complex and accurate, teams of developers start to look at how they could be used to aid in medical diagnosis. Here several examples related to use of CNNs specifically for Alzheimer's studies and other neurodegenerative disorders, shall be explored.

First is a paper titled "Alzheimer's disease classification based on combination of multi-model convolutional networks". This paper discusses the potential for designing a 3D CNN specifically for processing MRI data, building layers to extract features from given brain scans. It goes on to list where they intend to get the data to train the network from, the Alzheimer's Disease Neuroimaging Initiative (ADNI) [9]. The set for training the looked as was roughly 50/50 split confirmed Alzheimer's cases, and a control group [10]. This study is rather interesting, in that it is the kind of study that this major project proposal would ideally be if it were for a higher level than BSc. Their approach looks at directly handling the 3D data, which is the preferred way, for this project the data would be handled using a 2D image architecture but with multiple channels, essentially splitting each 3D image into an array of layers. The source of the data they used looks to be the best option [11] (out of a collection of various open access medical image repositories), though it would be interesting to see if a larger dataset could be used; using a larger dataset could help to make up for using multichannel 2D convolution, but it would also mean more data to pre-process and an increased training time. Their study would be interesting to revisit at MSc or PHD level.

Next is a study that looked at diagnosis of Alzheimer's at the early stage through use of deep learning. This study from 2014 brings up that at the time, previous methods used for diagnosis utilising machine/deep learning has reached a bottleneck of sorts, and so they looked at trying make a new architecture with aim of getting past said bottle neck. At the end of their experiments they found increased performance in all diagnosis datasets that they used [12].

A 2016 article published in the 'International Journal of Neural Systems', acknowledges the importance of computer aided diagnosis (CAD), and uses the remainder of the article to discuss what classification methods would be best to use in the diagnosis of Alzheimer's. Part of the basis of the article is the idea of using ensembles of networks to improve performance/accuracy at the cost of needing increased computing power [13]. For CAD ensembles of networks seem to be the preferred option compared to node dropout, as accuracy is the highest priority when a patient's health is on the line. For a repeat of this proposed project network ensembles would be a preferred option to use, but at this stage dropout would be more realistic especially with the training likely to be carried out on a single computer – a computer with hardware not specifically designed with network training in mind.

Last is a study looking at Parkinson's disease. This study used an arguably quite shallow architecture for their deep learning approach, with only 13 layers. This architecture was trained using electroencephalogram (EEG) data, as it aims to help diagnose Parkinson's at the early stages [14]. Their study looks to have had a fairly small dataset, 20 healthy subjects and 20 confirmed Parkinson's patients; a small dataset could be problematic as they claim to have achieved high levels of accuracy, high accuracy with small datasets could mean that the network has overfitted. While not related to Alzheimer's, the study still relates to a neurodegenerative condition (so relates to the interest of the author) and demonstrates a use case of CNNs in for CAD.

3 – Proposed experiment based on MRI datasets

An experiment is to be carried out to compare the effectiveness of a number of pretrained network architectures when trained with MRI data to try and diagnose Alzheimer's cases. The project shall be carried out using MATLAB, making use of the deep learning toolbox.

The project shall consist of:

- Finding a dataset to use
- Identifying which architectures to use for transfer learning (repurposing a pretrained network for your purposes and training with a new set of data)
- Preparing the data for training with said architectures
- Modifying the network layers of the pretrained networks if necessary (the input to be able to handle the multichannel data, and the fully connected + classification layers)
- Pre-processing the dataset to work with the pretrained networks (different architectures take different resolution images as input, so pre-processing should be done using functions and a datastore when the images are loaded into the programming environment, so as not to take up more hard drive space than absolutely necessary)
- Training the networks (all network training to either be done on the same computer, one after the other, or simultaneously on multiple computers with the same hardware)
- Recording results of training (training time for current training options on the given architecture, accuracy, etc)

At a minimum, one pretrained network should get retrained. Ideally three different architectures will be used to give a range of data to compare at the end. Given enough time multiple iterations could be trained by altering training options; to make sure that it is just the architectures being changed, the same training options should be used for each architecture – if training options are changed for one in an iteration, change it for all of them. At a stretch, an interesting goal would be to carry out

the same training multiple times on multiple computers with the same hardware to be able to build confidence in the data being collected.

4 – Research and development methods

The project follows a positivist research methodology, collecting scientific data and analysing the results. The project will also be following an inductive approach, with no prediction being made about what the outcomes of the experiments will be but observing the results with the aim of understanding more about the subject area.

The method of collecting data is through observation, recording the results from the training window in the MATLAB environment and the end of each network training.

There is a high level of risk associated with many aspects of the project, the biggest being the potential ethical and legal ramifications of the data that will be used for the network training. The dataset is medical data, which means that great care needs to be taken to follow the guidelines set out by the organisation who holds the rights to the data; one of the biggest parts of the being the protection of the identity of the patients the scans belong too – the scans hold no data to link back to an individual person, but at no time should an individual try to learn the identity of the patients. Any other person who may have access to the dataset – lecturers who may be monitoring the progress of the project – also need to follow the guidelines for access of the data. At no point will the data be shared with any other student, if another student wishes to use the data, they should follow the same path to gain permission to access the data.

5 – Resources and planning

The main resources that will be required are a computer for training the networks, and the data to train them with. The computer used will likely be that of the author, if at all possible though multiple university computers would be used so to be able to repeat training on computers with the same hardware; multiple instances of training the same architecture would be started at the same time so they would all be unavailable to use while the training is taking place (if the computers are used for other tasks during training, that would skew the results). MATLAB will be the programming language/development environment of choice; this is easily available as an academic license is available from the university.

The dataset being used will ideally be ADNI [9], this means permission needs to be given from them to use said data and they have their own usage policy to ensure data-protection. This has the greatest potential for risk, if permission is not given it might be difficult to find another relevant dataset.

Time should be spent over the summer before the next term starts to begin progress on the project. One of the key parts of this will be following the correct process to get permission to use the dataset. The other main plan is to spend this time learning more about network training, and about neuroanatomy and MRI scans, so that a strong level of knowledge is held once the next term starts.

A Gantt chart has been made with a rough timeline for the proposed project:

		Task Mode ▾	Task Name ▾	Duration ▾	Start ▾	Finish ▾	Predecessors ▾	Resource Names
1			Major Project	263 days	Sun 31/05/20	Tue 01/06/21		
2			Gain access to dataset	21 days	Sun 31/05/20	Fri 26/06/20		
3			Background research	89 days	Sun 31/05/20	Wed 30/09/20		
4			Modify architectures	14 days	Thu 01/10/20	Tue 20/10/20	3	
5			Write data-processing functions	28 days	Thu 01/10/20	Mon 09/11/20	2,3	
6			Architecture 1 training 1	14 days	Tue 10/11/20	Fri 27/11/20	2,4,5	
7			Architecture 2 training 1	14 days	Mon 30/11/20	Thu 17/12/20	6	
8			Architecture 3 training 1	14 days	Fri 18/12/20	Wed 06/01/21	7	
9			Architecture 1 training 2	14 days	Thu 07/01/21	Tue 26/01/21	8	
10			Architecture 2 training 2	14 days	Wed 27/01/21	Mon 15/02/21	9	
11			Architecture 3 training 2	14 days	Tue 16/02/21	Fri 05/03/21	10	
12			Write up findings	57 days	Mon 08/03/21	Tue 25/05/21	11	

Figure 1 Gantt chart screenshot 1

Remaining screenshots available at the following links:

- <https://i.imgur.com/gyzz6Ol.png>
- <https://i.imgur.com/pwpPGz2.png>
- <https://i.imgur.com/reAJfyC.png>

References

- [1] ImageNet (Stanford University), "ImageNet Large Scale Visual Recognition Challenge (ILSVRC)," ImageNet, 2017. [Online]. Available: <http://www.image-net.org/challenges/LSVRC/>. [Accessed 27 April 2020].
- [2] Y. LeCun, L. Bottuo, Y. Bengio and P. Haffner, "Gradient-Based Recognition Learning Applied to Document Recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 1-46, 1998.
- [3] s. Das, "CNN Architectures: LeNet, AlexNet, VGG, GoogLeNet, ResNet and more...," 16 November 2017. [Online]. Available: <https://medium.com/analytics-vidhya/cnns-architectures-lenet-alexnet-vgg-googlenet-resnet-and-more-666091488df5>. [Accessed 27 April 2020].
- [4] A. Krizhevsky, I. Sutskever and G. Hinton, "ImageNet Classification with Deep Convolutional," in *Neural Information Processing Systems*, Tahoe (Nevada), 2012.
- [5] J. Brownlee, "A Gentle Introduction to Dropout for Regularizing Deep Neural Networks," 6 August 2019. [Online]. Available: <https://machinelearningmastery.com/dropout-for-regularizing-deep-neural-networks/>. [Accessed 27 April 2020].
- [6] M. Zeiler and F. R, "Visualizing and Understanding Convolutional Networks," in *European Conference on Computer Vision*, Zurich, 2014.
- [7] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke and A. Rabinovich, "Going Deeper with Convolutions," in *IEEE Conference on Computer Vision and Pattern Recognition*, Boston, 2015.
- [8] K. He, X. Zhang, S. Ren and J. Sun, "Deep Residual Learning for Image Recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, 2016.
- [9] Alzheimer's Disease Neuroimaging Initiative, "Alzheimer's Disease Neuroimaging Initiative," 2017. [Online]. Available: <http://adni.loni.usc.edu/>. [Accessed 28 April 2020].
- [10] F. Li, C. D and M. Liu, "Alzheimer's disease classification based on combination of multi-model convolutional networks," in *IEEE International Conference on Imaging Systems and Techniques*, Beijing, 2017.
- [11] S. Aylward, "Open-Access Medical Image Repositories," [Online]. Available: <http://www.aylward.org/notes/open-access-medical-image-repositories>. [Accessed 28 April 2020].
- [12] S. Liu, S. Liu, W. Cai, S. Pujol, R. Kikinis and D. Feng, "Early diagnosis of Alzheimer's disease with deep learning," in *IEEE 11th International Symposium on Biomedical Imaging*, Beijing, 2014.
- [13] A. Ortiz, J. Munilla, J. Górriz and J. Ramírez, "Ensembles of Deep Learning Architectures for the Early Diagnosis of the Alzheimer's Disease," *International Journal of Neural Systems*, vol. 26, no. 7, 2016.

- [14] S. Oh, Y. Hagiwara, U. Raghavandra, R. Yuvaraj, N. Arunkumar, M. Murugappan and R. Acharya, "A deep learning approach for Parkinson's disease diagnosis from EEG signals," *Neural Computing and Applications*, 2018.