Proceedings of the 2016 International Conference on Machine Learning and Cybernetics, Jeju, South Korea, 10-13 July, 2016

LEPROSY LESION RECOGNITION USING CONVOLUTIONAL NEURAL NETWORKS

HARJATIN SINGH BAWEJA[1] *, TANVIR PARHAR[2] *

[1]School Of Computer Science, Student of Bachelors in Computer Science, Vellore Institute Of Technology, Vellore, 632014, India

[2] School Of Electrical Engineering, Student of Bachelors in Electronics and Instrumentation, Vellore Institute Of Technology, Vellore, 632014, India

E-MAIL: harjatin@gmail.com, tanparhar@gmail.com

Abstract:

Leprosy, also known as Hansen's disease, is a debilitating and chronic bacterial infection. As per World Health Organization's report, there were 189,000 chronic cases of Leprosy in 2012 with 230,000 new diagnoses. Although curable at later stages, an early diagnosis prevents nerve involvement and the disabilities it incurs. The authors henceforth propose a Convolutional Neural Network based architecture for Leprosy lesion recognition. To train the network, authors use DermnetNz datasets along with web scraped images to achieve a best accuracy of 91.6% on a dataset split into 60% of training images, 20% of images are used for cross validation and 20% for testing.

Keywords:

Convolutional Neural Networks; Computer Vision; Leprosy; Hansen's disease; Artificial Neural Networks; Tensor Flow

1. Introduction

As suggested by World Health Organization's report in 2013, most cases of Hansen's disease stem out from high prevalence areas such as South Asia, Brazil, some regions in Africa and Western Pacific. In fact, as few as 14 countries account for 95% of the total cases in the world. Within these countries, India accounts for the highest i.e. around 60% of the total cases, followed by Brazil with 14% and Indonesia with 8%.

Although curable, a delay in diagnosis of Hansen's disease can lead to several irreparable damages such as disfigurement, blindness, glaucoma, kidney failure, erectile dysfunction, infertility and inability to use hands and feet. However, studies conducted by Patricia et.al. on 506 patients suggest a mean age of diagnosis as 41 with an average delay of just over 2 years [1]. Although there is no formal definition on what constitutes an acceptable delay for detection after the first symptom shows, Muller [2] suggests that a delay longer than 6 months is deleterious to the patient and to the economy. In case of Neurotic Leprosy, a delay in diagnosis of over 5 years could mean a 15-18 times increase in risk of deformity. Whereas paralytic deformities are rare for a delay period of less than 1 year. [3]

In recent years, with the exponential increase in computing power, neural networks have made significant performance improvements in many fields of artificial intelligence. Convolutional Neural networks presents a unique architecture that allows training on high dimensional data such as images while using reasonable computing resources. It does so by combining three architectural ideas to ensure some degree of shift and distortion invariance: local receptive field, shared weights, and sometimes spatial or temporal subsampling.

Over the years, there have been several variations of convolutional neural network architectures used for object classification [4],[5],[6],[7],[8]. However, the state-of-the-art architecture at present is Google's inception-v3 [9] which has reported 3.5% top-5 error and 17.3% top-1 error on the ImageNet Large Scale Visual Recognition Challenge [10].

Hence, the authors propose to use the factorized computation of convolution operations as seen in state of the art Google's inception-v3 architecture for Leprosy Lesion identification. The network is trained with a dataset of 150 RGB 32x32 images. The ratio of testing: validation: training is 1:1:3. Each group consists of 50% of positive and negative samples. Negative image dataset consists of images from skin diseases like skin cancer, sores and mild burns. Almost all vital parts of the body have been covered in the dataset.

2. Related Work

Surprisingly not much work has been done in the field of image based Leprosy Lesion detection. Contact based methods that measure the exponential delay in the passing of action potential across a diseased nerve have been implemented [11]. Machine learning based approaches have been implemented to detect the subtle variations in

978-1-5090-0390-7/16/$31.00 ©2016 IEEE

myo-electric signals in the Lesion afflicted parts of the body [12]. The problem with such techniques is that they are either invasive or contact based which often require complex and expensive machinery.

Computer vision based techniques have been implemented for detection of skin cancer at early stages. However, such techniques rely on manual feature extraction and labelling, thus leading to a requirement of human-resources. Machine learning algorithms like SVM have been used to classify the features into classes of skin diseases.[13].

3. Approach

The approach suggested by the authors in this paper uses state-of-the-art computer vision technique inspired by Google's inception-v3 model used in the most recent ILSVRC-2015 challenge with images of both Leprosy and non-Leprosy skin lesions. Owing to properties such as factorized convolutions and aggressive regularization, the computational efficiency of this architecture would enable easy scaling of the general purpose system to various mobile platforms.

To better sample the visual domain, images were randomly shuffled in the training dataset and various image operations such as random translations, skew transformations, scale transformations, blurring effects and rotations were used on both training and testing subdivisions. This makes the system robust to changes in orientations, scale and to some extent camera properties also. These properties make the system deployable on mobile platforms as well.

3.1. Visual features of Leprosy Lesion

A diagnosis by a dermatologist would include a visual pre-screening of the skin lesion along with a filament test on specific areas or a simple feather test to check for loss of sensation on the lesion.

However, a leprosy lesion is characterized by certain distinct features which are used by both, an experienced dermatologist as well as a trained ensemble of convolutional neural network architecture to differentiate it from other similar looking skin patches. These features would include hypo-pigmented skin lesion with no hair growth on it and different textures describing remarkable features of the lesion, as shown in Figure 1. It is easy to confuse the lesion with other skin problems like dermatitis (skin allergies), ring worm infection, syphilis or psoriasis. Thus making the recognition of the lesion a hard task for an untrained eye.

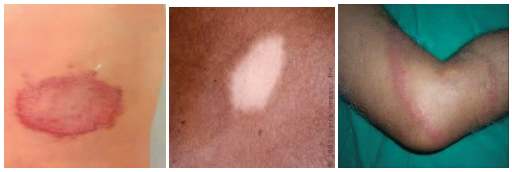

FIGURE 1. Visual Features of Leprosy Skin Lesion. From Left to right A) Difference in Skin Texture B) Hypo-Pigmentation C) Retarded Hair Growth

3.2. Datasets

The image dataset consists of a batch of equally distributed positive and negative examples of both leprosy and non-leprosy lesions. It was collected from a popular skin vision dataset available on www.dermnetnz.org and web scraping of google image search results and manually cleaning them. Figure 2 shows some sample images from the dataset.

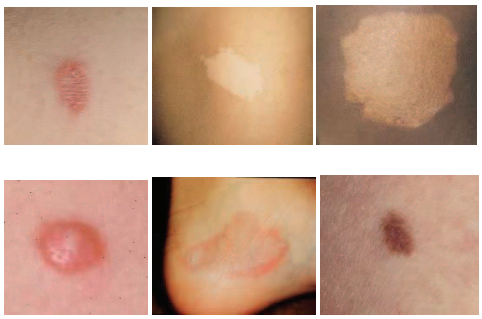

FIGURE 2. Top three images represent positive training examples and bottom three represent negative training examples

Since the training dataset consisted of only 120 training and validation examples, the image samples were randomly geometrically transformed within each training cycle in order to maximize the effectiveness of the collected images .

3.3. CNN Architecture

For purpose of experimentation, the authors have used Google's open source software library, TensorFlow, for Machine Intelligence. Specifically, TensorFlow's factorized implementation of state-of-the-art architecture for Convolutional neural networks was used. As shown in Figure 3, the network takes in a batch of randomly selected twenty images during each training step. The images are convolved with a number of filters. Each filter learns to detect a specific

feature in the image which could be a gradient orientation, stroke, curvature, etc. A number of such features at different image locations combine to describe the image in feature space. The weights of each feature map are shared across the entire image through the RGB color-space to look for a particular feature in the image irrespective of its location. The output of this layer is passed through the Rectified Linear Unit (popularly known as ReLu) which forces the output to be between zero and one. This function has an additional advantage over its counterparts such as softmax function of being able to effectively avoid the vanishing gradient problem in the back propagation step thus leading to faster convergence while training of the network. A smooth approximation of this function, known as softplus function, is represented by equation 1.

$$f(x) = \ln(1 + e^x) \tag{1}$$

The output of this layer is max-pooled over a window of 2x2 so as to reduce the dimensionality of the inputs for the next convolutional layer while simultaneously achieving small translational equivarience. The max-pooled layers are then convolved with another set of filters, repeating the same procedure as in first convolutional layer but learning more subtle features. The convolved output is then passed through another ReLu function.

The outputs of the second convolutional layer are connected to fully connected neurons which are in turn connected to two neurons, each giving a probability of the image being a positive example or a negative example respectively through a softmax function represented in equation 2. In the equation, Z represents a K dimensional vector and j goes from 1 to k.

$$\sigma(\mathbf{z})_j = \frac{e^{z_j}}{\sum_{k=1}^{K} e^{z_k}} \tag{2}$$

4. Experiments

4.3. Optimization Strategies

Three flavors of gradient descent based optimization strategies were used: Batch Gradient Descent [14], Momentum Optimizer [15] and Adaptive Moment Optimizer (ADAM) [16]. Cross entropy was used as the cost function [17].

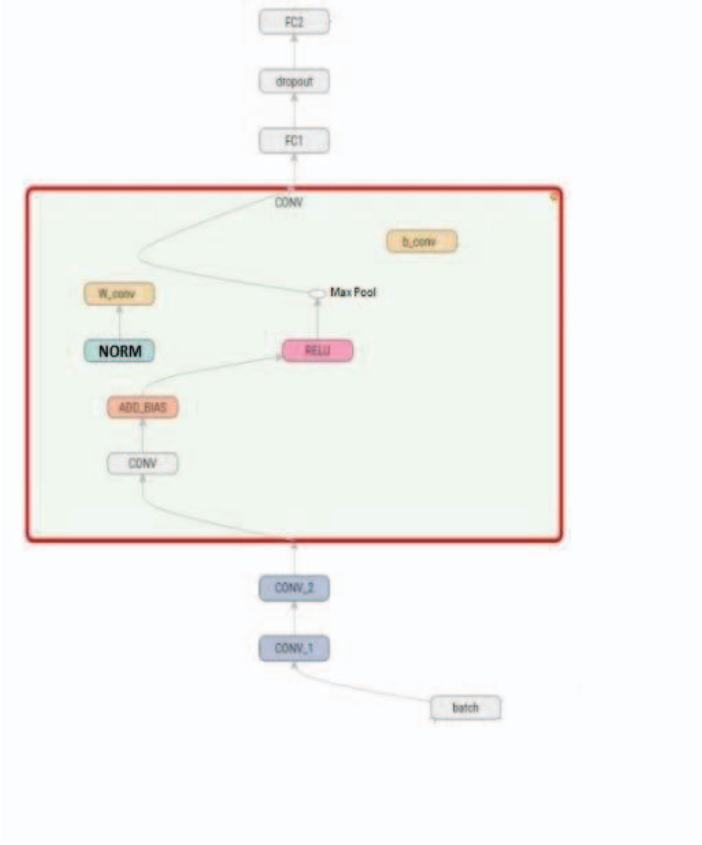

FIGURE 3. Convolutional Neural Network Architecture

As seen from Figure 4, a naïve Batch Gradient Descent based optimizer performs worst. It has oscillations that result in late convergence of the network. To minimize the oscillations of the Batch Gradient Descent, Optimizer Momentum based optimizer is used. The Momentum optimizer takes into account the gradients from previous iterations. This reduces the oscillations and the number of iterations to converge as well. However, a drawback of Momentum optimizer is that it does not use learning rate adaptation with respect to each parameter of the cost function. Hence, an Adaptive Moment Estimation (ADAM) optimizer is implemented which shows better convergence results. It computes the adaptive learning rates for each parameter as well as keeps an exponentially decaying average of past gradients, similar to momentum.

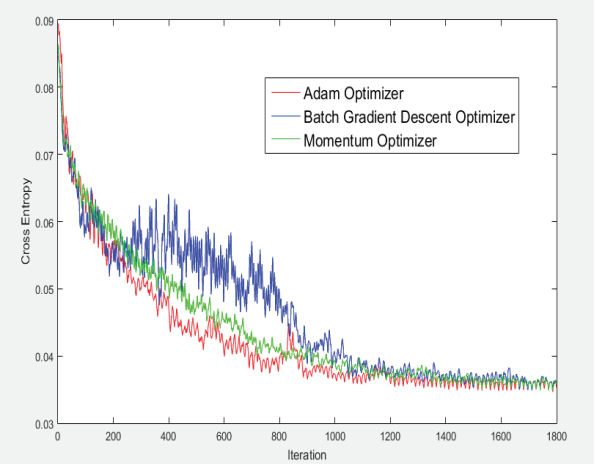

FIGURE 4. Cross Entropy vs Iteration for various optimization techniques

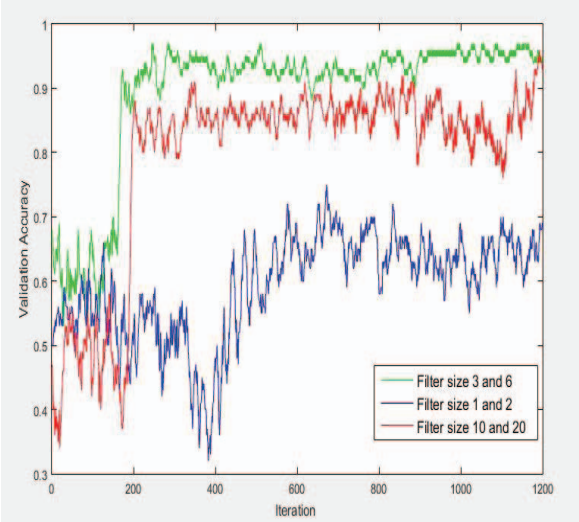

FIGURE 5. Validation accuracy vs Iteration for different number of filters in convolutional layer 1 and layer 2

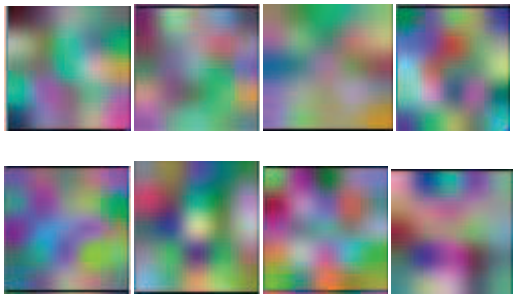

FIGURE 6. RGB descriptions of learned features of layer1 and layer 2

The trade-off for better convergence results is the time taken in each iteration. Since Adam optimizer requires most computations per training step, it takes the maximum time of 1791.896 sec, followed by1789.9 sec for Momentum Optimizer and 1773.06 sec for Batch Gradient Descent Optimizer. This was computed for 20,000 training iterations on an octa-core Intel i-7 3770k processor. Since ADAM Optimizer gives the best convergence results, the results obtained by a network trained by ADAM Optimizer are discussed in the trailing sections.

4.4. Classification Results

Figure 5 shows the Validation accuracy as a function of iteration number for different depths of feature maps for Convolutional layer 1 and Convolutional layer 2 respectively.

As inferred from the graph in figure 4, the optimal feature map depth for convolutional layer1 is 3 and for convolutional layer 2 is 6(Green). Lesser depth in both the layers (Blue) decreases the accuracy significantly, as the output feature map fails to sufficiently capture key spatial features of the image. It is also seen that if the number of filters for both the layers is increased beyond an optimal number (Red), the network over-fits on the training set while failing to generalize on validation dataset and thus resulting in higher accuracies for training dataset but lower accuracies for validation set. Average accuracy for used architecture is 91.6%, whereas average accuracy for network with lesser number of filters is 69.4% and for network with higher number of features is 86.1%. Some of the learned filters for optimal network architecture are shown in interpretable form in figure 6, after normalization in 8 bit RGB space.

5. Conclusions

This paper re-engineers the problem of Leprosy lesion detection into an image recognition problem and implements the state-of the art Convolutional Neural Network architecture to solve it. The optimal number of filters and optimizers to train the network, for maximizing the accuracy on a limited dataset is found empirically. The final architecture uses an ADAM optimizer since obtaining highest accuracy is of most priority and training time is not a constraint. It can be seen in Figure 6 that the network learns to detect more spread out features in the first layers, as compared to more well-rounded features in the second layer filters. The final architecture achieves an accuracy of 91.6%

Current methods require dedicated hardware and software systems and trained professionals for diagnosis of Hansen's disease. Such facilities are generally not available in third world countries leading to delayed diagnosis, thus irrepealably damages. The limited vision based techniques

suggested that inputting features that require human expertise for proper sampling. These inadequacies along with the stigma attached with the disease often lead patients into ignorance followed by disastrous consequences. This research lays the ground work for development of a mobile image analysis system for diagnosis of Leprosy Lesion in domestic environments, thus making it free of the above mentioned shortcomings. The hope of this research is to provide clinicians as well as domestic users with an easy to use non-invasive vision based tool to aid early diagnosis of skin diseases.

Acknowledgements

The authors would like to acknowledge VIT University for providing with the necessary resources for completing the research.

References

[1] Deps, Patricia D., et al. "Delay in the diagnosis of leprosy in the Metropolitan Region of Vitoria, Brazil." Leprosy review 77.1 (2006): 41-47.

[2] Muller, P., et al. "Lèpre en Guadeloupe: maladie en déclin, délaidiagnostique en hausse." Annales de dermatologieet de vénéréologie. Vol. 130. No. 6-7. Masson, 2003.

[3] Kumar, Anil, ANITA GIRDHAR, and B. K. Girdhar. "Nerve thickening in leprosy patients and risk of paralytic deformities: a® eld based study in Agra, India." Lepr Rev 75 (2004): 135-142.

[4] Krizhevsky, Alex, IlyaSutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." Advances in neural information processing systems. 2012.

[5] Szegedy, Christian, et al. "Going deeper with convolutions." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2015.

[6] LeCun, Yann, et al. "Gradient-based learning applied to document recognition." Proceedings of the IEEE 86.11 (1998): 2278-2324.

[7] He, Kaiming, et al. "Deep Residual Learning for Image Recognition." arXiv preprint arXiv:1512.03385 (2015).

[8] Simonyan, Karen, and Andrew Zisserman. "Very deep convolutional networks for large-scale image recognition." arXiv preprint arXiv:1409.1556(2014).

[9] Szegedy, Christian, et al. "Rethinking the Inception Architecture for Computer Vision." arXiv preprint arXiv:1512.00567 (2015).

[10] Olga Russakovsky*, Jia Deng*, Hao Su, Jonathan Krause, SanjeevSatheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, AdityaKhosla, Michael Bernstein, Alexander C. Berg and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. IJCV, 2015. paper

[11] Vivekanandan, S., D. S. Emmanuel, and RichaKumari. "Propagation of Action Potential For Hansen's Disease Affected Nerve Model Uisng Fitzhugh Nagumo Like Excitation." Journal of Theoretical and Applied Information Technology 49.2 (2013).

[12] Vivekanandan, S., et al. "Classification of Myoelectric Signals to Diagnose Hansen's Disease.

[13] Maglogiannis, Ilias, and Charalampos N. Doukas. "Overview of advanced computer vision systems for skin lesions characterization." Information Technology in Biomedicine, IEEE Transactions on 13.5 (2009): 721-733.

[14] LeCun, Yann A., et al. "Efficient backprop." Neural networks: Tricks of the trade. Springer Berlin Heidelberg, 2012. 9-48.

[15] Qian, Ning. "On the momentum term in gradient descent learning algorithms." Neural networks 12.1 (1999): 145-151.

[16] Kingma, Diederik, and Jimmy Ba. "Adam: A method for stochastic optimization." arXiv preprint arXiv:1412.6980 (2014).

[17] Nasr, George E., E. A. Badr, and C. Joun. "Cross Entropy Error Function in Neural Networks: Forecasting Gasoline Demand." FLAIRS Conference. 2002.