Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 90 (2016) 200 – 205

Procedia Computer Science

International Conference On Medical Imaging Understanding and Analysis 2016, MIUA 2016, 6-8 July 2016, Loughborough, UK

Convolutional Neural Networks for Diabetic Retinopathy

Harry Pratt[a,*], Frans Coenen[b], Deborah M Broadbent[c], Simon P Harding[a,c], Yalin Zheng[a,c]

[a]*Department of Eye and Vision Science, Institute of Ageing and Chronic Disease, University of Liverpool, Apex Building, 6 West Derby Street, Liverpool L7 9TX, United Kingdom*

[b]*Department of Computer Science, University of Liverpool, Ashton Street, Liverpool L69 3BX, United Kingdom*

[c]*Royal Liverpool University Hospital, St. Paul's Eye Unit, Prescot Street, Liverpool L7 8XP, United Kingdom*

Abstract

The diagnosis of diabetic retinopathy (DR) through colour fundus images requires experienced clinicians to identify the presence and significance of many small features which, along with a complex grading system, makes this a difficult and time consuming task. In this paper, we propose a CNN approach to diagnosing DR from digital fundus images and accurately classifying its severity. We develop a network with CNN architecture and data augmentation which can identify the intricate features involved in the classification task such as micro-aneurysms, exudate and haemorrhages on the retina and consequently provide a diagnosis automatically and without user input. We train this network using a high-end graphics processor unit (GPU) on the publicly available Kaggle dataset and demonstrate impressive results, particularly for a high-level classification task. On the data set of 80,000 images used our proposed CNN achieves a sensitivity of 95% and an accuracy of 75% on 5,000 validation images.

Peer-review under responsibility of the Organizing Committee of MIUA 2016

Keywords: Deep Learning, Convolutional Neural Networks, Diabetic Retinopathy, Image Classification, Diabetes

1. Introduction

Diabetic Retinopathy (DR) is one of the major causes of blindness in the western world[1 2]. Increasing life expectancy, indulgent lifestyles and other contributing factors mean the number of people with diabetes is projected to continue rising[3]. Regular screening of diabetic patients for DR has been shown to be a cost-effective and important aspect of their care[4]. The accuracy and timing of this care is of significant importance to both the cost and effectiveness of treatment. If detected early enough, effective treatment of DR is available, making this a vital process[5].

Classification of DR involves the weighting of numerous features and the location of such features[6]. This is highly time consuming for clinicians. Computers are able to obtain much quicker classifications once trained, giving the ability to aid clinicians in real-time classification. The efficacy of automated grading for DR has been an active area

* Harry Pratt. Tel.: +447428611330
E-mail address: sghpratt@liverpool.ac.uk

Peer-review under responsibility of the Organizing Committee of MIUA 2016
doi:10.1016/j.procs.2016.07.014

of research in computer imaging with encouraging conclusions[7 8]. Significant work has been done on detecting the features of DR using automated methods such as support vector machines and k-NN classifiers[9]. The majority of these classification techniques are on two class classification for DR or no DR.

Convolutional Neural Networks (CNNs), a branch of deep learning, have an impressive record for applications in image analysis and interpretation, including medical imaging. Network architectures designed to work with image data were routinely built already in 1970s[10] with useful applications and surpassed other approaches to challenging tasks like handwritten character recognition[11]. However, it wasn't until several breakthroughs in neural networks such as the implementation of dropout[12], rectified linear units[13] and the accompanying increase in computing power through graphical processor units (GPUs) that they became viable for more complex image recognition problems. Presently, large CNNs are used to successfully tackle highly complex image recognition tasks with many object classes to an impressive standard. CNNs are used in many current state-of-the-art image classification tasks such as the annual ImageNet and COCO challenges[14 15].

Two main issues exist within automated grading and particularly CNNs. One is achieving a desirable offset in sensitivity (patients correctly identified as having DR) and specificity (patients correctly identified as not having DR). This is significantly harder for national criteria which is a five class problem in to normal, mild DR, moderate DR, severe DR, and proliferative DR classes. Furthermore, overfitting is a major issue in neural networks. Skewed datasets cause the network to over-fit to the class most prominent in the dataset. Large datasets are often massively skewed. In the dataset, we used less than three percent of images came from the 4th and 5th class, meaning changes had to be made in our network to ensure it could still learn the features of these images.

In this paper, we introduce a deep learning-based CNN method for the problem of classifying DR in fundus imagery. This is a medical imaging task with increasing diagnostic relevance, discussed earlier, and one that has been subject to many studies in the past. As far as we are aware, this is the first paper discussing the five class classification of DR using a CNN approach. Several new methods are introduced to adapt the CNN to our large dataset. We then analyse the performance and dissect the capabilities of our network.

The remainder of this paper is organised as follows. Section 2 presents an overview of related work, section 3 describes the architecture of the CNN and the training methods used in this work, section 4 presents the results from our experiments, section 5 concludes the paper with discussion on the results and future work.

2. Related Work

Extensive research has been carried out on methods for a binary classification of DR with encouraging results. Gardner et al used Neural Networks and pixel intensity values to achieve sensitivity and specificity results of 88.4% and 83.5% respectively for yes or no classification of DR[16]. They used a small dataset of around 200 images and split each image in to patches and then required a clinician to classify the patches for features before SVM implementation.

Neural Networks have also been used in three-class classification of DR. Nayak et al[17] used features such as the area of exudates and the area of blood vessels together with texture parameters. Features are entered into the neural network to classify images into normal, non-proliferative retinopathy and proliferative retinopathy. The neural network used these features as input for classification. The detection results were validated by comparing with grading from expert ophthalmologists. They demonstrated a classification accuracy of 93%, sensitivity of 90% and specificity of 100%. This was carried out on a dataset of 140 images and feature extraction was required on all images in both training and testing which can be time consuming.

The vast majority of research on the five-class classification that has been carried out has used support vector machines (SVMs). Acharya et al[18] have created an automated method for identifying the five classes. Features, which are extracted from the raw data using a higher-order spectra method, are fed in to the SVM classifier and capture the