Predicting Blood Pressure from the Retina using Deep Learning

Dr. Alexander Button

Solving biological problems that require math

March 4, 2022

Retinal Image and the Cardiovascular System

The cardiovascular system is connected throughout the body. The eyes provide an "easy access" point to examine blood vessels directly.

In an example from 2014, a group from Google Research trained a model to predict a series of cardiovascular risk factors from retinal fundus images [1].

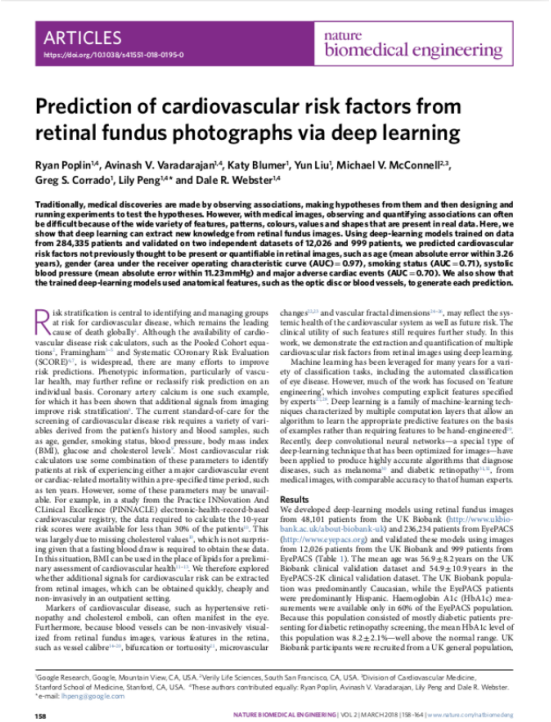

ARTICLES
https://doi.org/10.1038/s41551-018-0195-0

nature biomedical engineering

Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning

Ryan Poplin[1,4], Avinash V. Varadarajan[1,4], Katy Blumer[1], Yun Liu[1], Michael V. McConnell[2,3], Greg S. Corrado[1], Lily Peng[1,4]* and Dale R. Webster[1,4]

Traditionally, medical discoveries are made by observing associations, making hypotheses from them and then designing and running experiments to test the hypotheses. However, with medical images, observing and quantifying associations can often be difficult because of the wide variety of features, patterns, colours, values and shapes that are present in real data. Here, we show that deep learning can extract new knowledge from retinal fundus images. Using deep-learning models trained on data from 284,335 patients and validated on two independent datasets of 12,026 and 999 patients, we predicted cardiovascular risk factors not previously thought to be present or quantifiable in retinal images, such as age (mean absolute error within 3.26 years), gender (area under the receiver operating characteristic curve (AUC) = 0.97), smoking status (AUC = 0.71), systolic blood pressure (mean absolute error within 11.23 mmHg) and major adverse cardiac events (AUC = 0.70). We also show that the trained deep-learning models used anatomical features, such as the optic disc or blood vessels, to generate each prediction.

Figure 1: 2014 paper [1]

Prediction from Images

In order to make these predictions, the team used what is known as a convolutional neural network (CNN).

A CNN is a neural network architecture inspired from the visual cortex of the brain. This method takes a two-dimensional image as an input and combines the various pixel values in order to make a prediction.

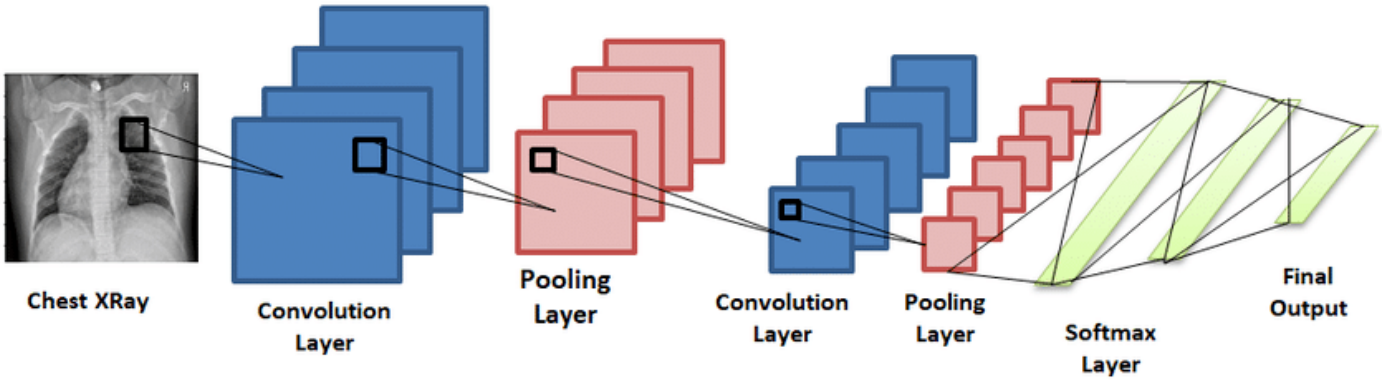

Figure 2: Overview of a convolutional neural network [2]

Retinal Image and the Cardiovascular System

Using their trained CNN, the team was able to predict various risk factors relevant for cardiovascular disorders.

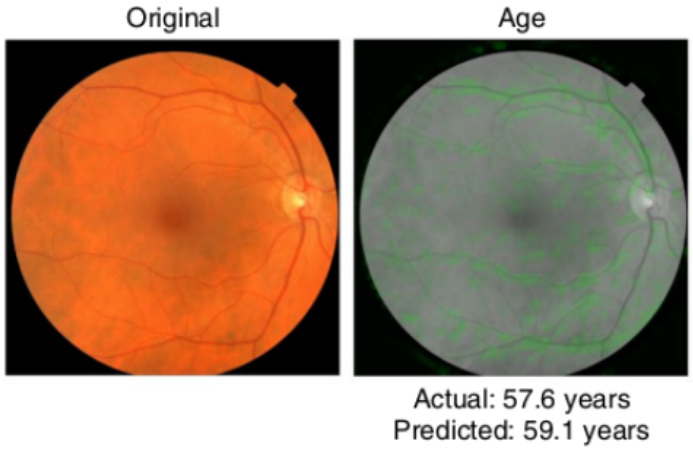

Figure 3: Raw fundus image (left) and the activation highlighted (right) for age prediction

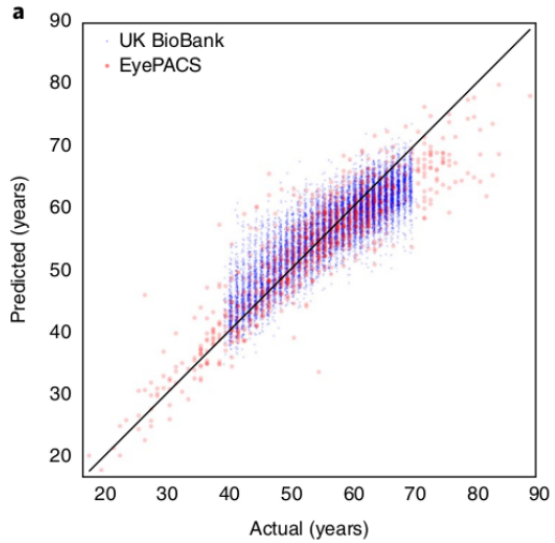

Figure 4: Comparison of predicted and actual age in years.

Aims of the Project

Aim 1

Training a model to predict different cardiovascular disorders directly from retinal images.

Aim 2

Extract the hidden layers of the model and determine how they contribute to the disease prediction.

Aim 3

Investigate the clinical and genetic significance of the extracted features.

Model Architecture

For this problem, we chose to use the **DenseNet** architecture. The DenseNet is a convolutional neural network that consists of a series of inter-connected **denseblock modules**.

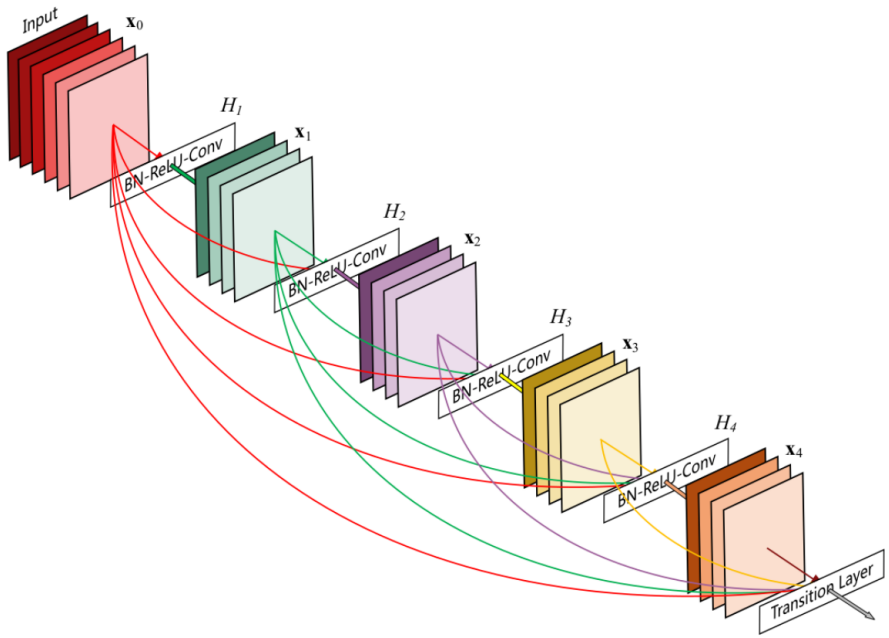

Figure 5: Overview of how the DenseNet works [3]

Input data

For our input data, we obtained ∼60,000 retinal images from the UKBioBank, as well as the subject's corresponding genomes.

As an initial preprocessing set, we converted these images into their vessel segmented forms.

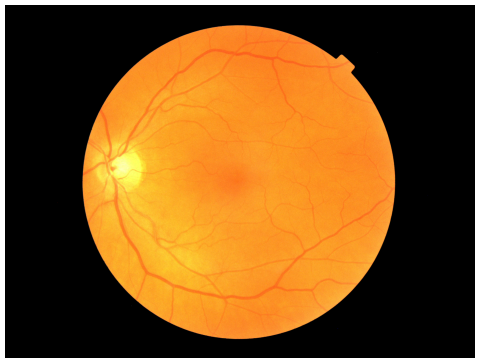

Figure 6: Raw retinal fundus image

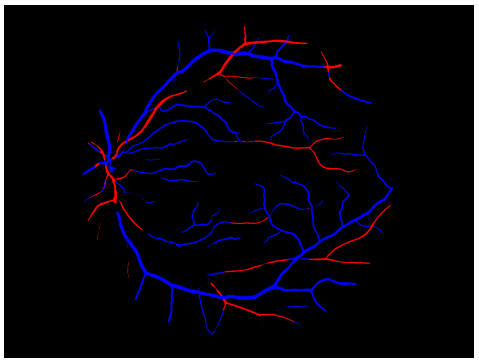

Figure 7: Vessel segmented image

Target Disorder

The key trait that we have been investigating has been **hypertension**. We predict hypertension categorically, that is, either as 0 (control) or 1 (case)

In our study we define hypertension as:

Control

DBP < 80 and SBP < 120

Case

DBP > 90 or SBP > 140

Prediction Results

We have successfully trained out model to predict hypertension and obtained a high predictive accuracy.

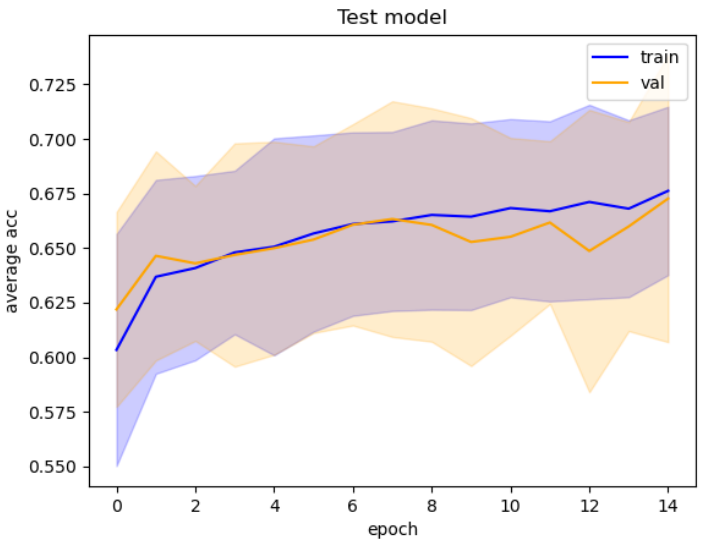

Figure 8: Training curve comparing the predictive accuracy between the training and validation data

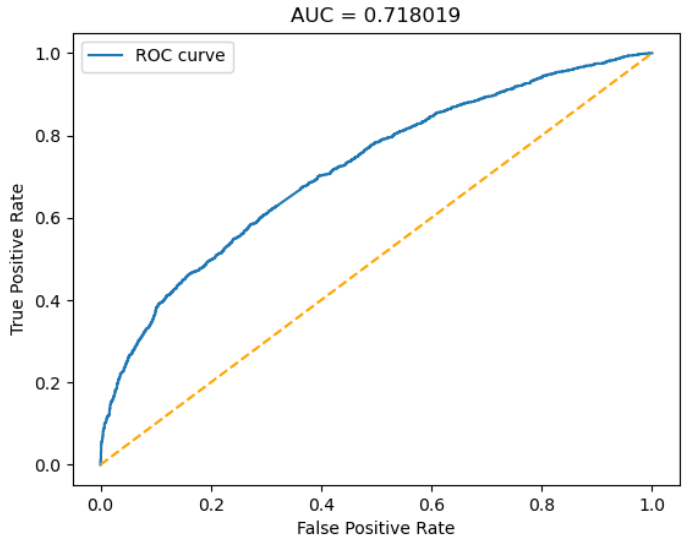

Figure 9: Receiver operator curve for predicting the hypertensive label (AUC=0.72)

Activation Visualization

To further understand which features of the image were contributing to the prediction, we generated **class activation maximized images**, that is, images in which the pixels are coloured according to their contribution to the output.

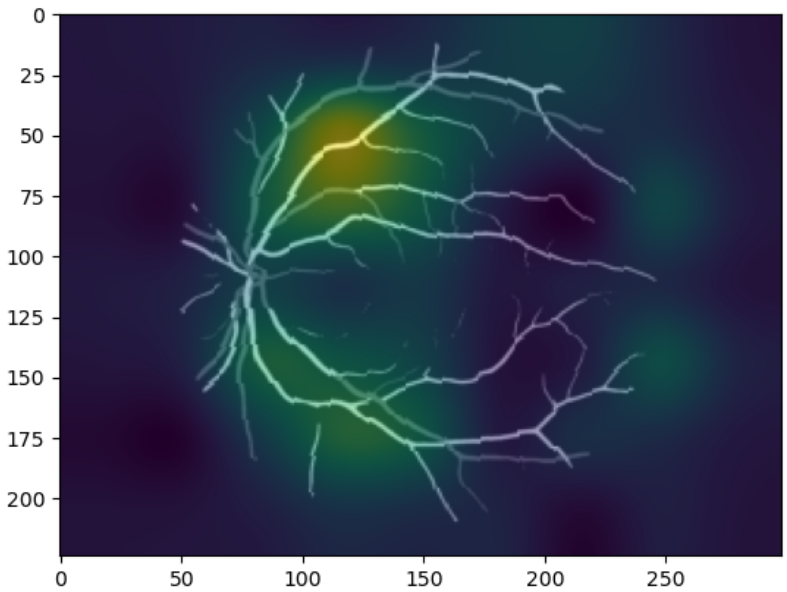

Figure 10: Vessel segmented image highlighted by the class activation

Genetic Analysis

We also performed a series of **genome-wide association studies (GWAS)** on a subset of the output layer neurons in order to determine their genetic significance.

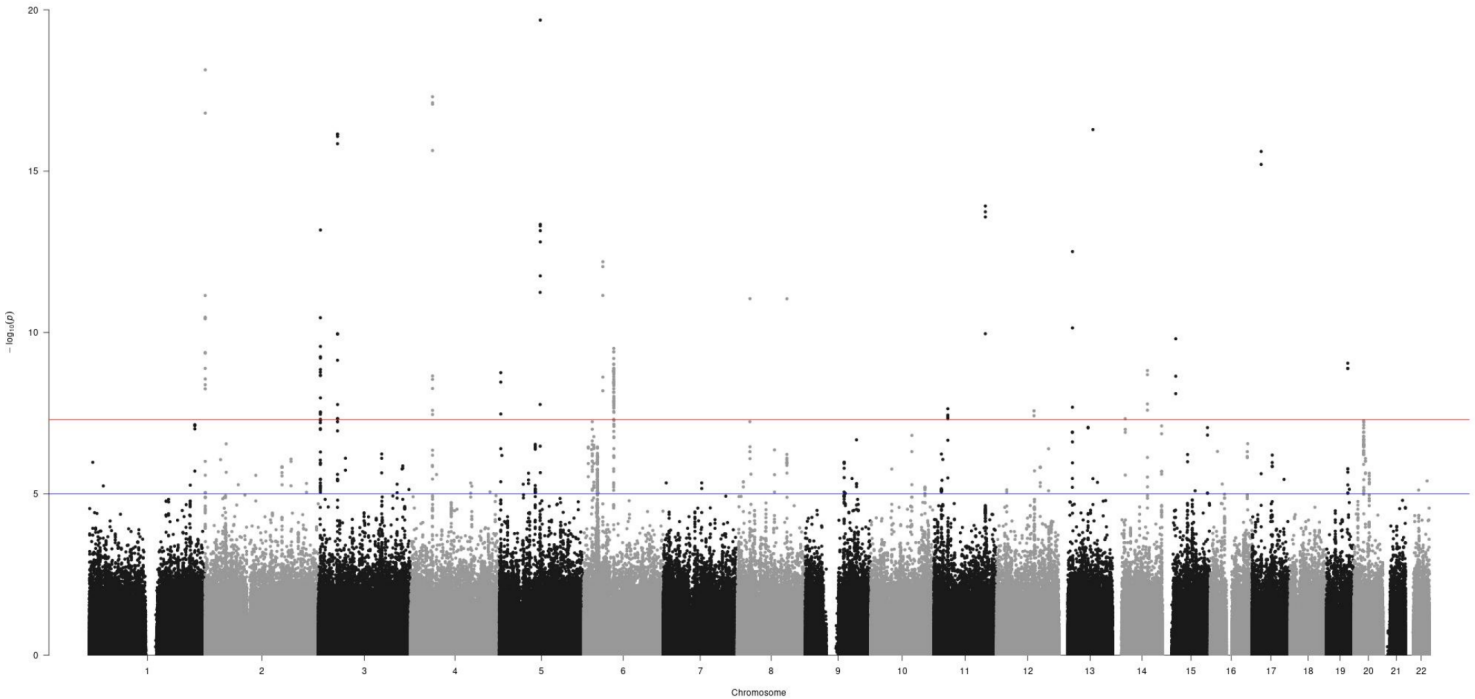

Figure 11: Manhatten plot for the activation values of a selected output layer neuron

Proposed Project

Project Proposal:

- The aim of the project would be to apply the same pipeline to different disease categories e.g. angina, diabetes, stroke, heart attack event.
- The method could be applied both on the raw images as well as the vessel segmented images
- One could also investigate in predicting continuous traits such as age, diastolic pressure, systolic Blood Pressure.

[1] Ryan Poplin, Avinash V. Varadarajan, Katy Blumer, Yun Liu, Michael V. McConnell, Greg S. Corrado, Lily Peng, and Dale R. Webster.
Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning.
Nature Biomedical Engineering, 2(3):158–164, 2018.

[2] Mohammad Farukh Hashmi, Satyarth Katiyar, Avinash G. Keskar, Neeraj Dhanraj Bokde, and Zong Woo Geem.
Efficient pneumonia detection in chest xray images using deep transfer learning.
Diagnostics, 10(6):1–23, 2020.

[3] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger.
Densely connected convolutional networks.
Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, 2017-Janua:2261–2269, 2017.