

Introduction

PROBLEM

Diabetic retinopathy (DR) is the most frequently occurring complication of diabetes mellitus. Over an extended period, high blood sugar levels can damage the eye's blood vessels and cause them to leak fluid or bleed, causing the formation of lesions in the retina¹. The severity of diabetic retinopathy falls on a range from mild to severe. In mild cases, there may be no symptoms or minor vision problems. If left untreated, severe cases may result in extreme vision loss or complete blindness. The exponential increase in the number of diabetics around the world has led to a corresponding increase in the number of diabetic retinopathy cases. The prolific growth of the health condition has created an increasing demand for qualified ophthalmologists to treat these cases. An automated DR detection system would significantly reduce the burden that is currently being placed on ophthalmologists.

APPROACH

The proposed approach that we used for the DR multi-class classification problem was composed of a network with an ensemble of feature extractors (Fig. 1) and an MLP classifier. The features outputted by each CNN are one-dimensionally concatenated into an array that is passed through an MLP classifier (Fig. 2).

Objective: Perform exploratory analyses of different classifier head designs to optimize the Diabetic Retinopathy classification model

Materials and Methods

DATA

*EyePACS Dataset*² is the Kaggle Diabetic Retinopathy dataset consisting of 88,702 color fundus images, including 35,126 samples for training and 53,576 samples for testing. An image of the left eye and the right eye were collected for each subject.

TECHNIQUES

Deep and Shallow Neural Network (DaSNN)

Feature extraction using deep learning:

A high-quality feature extractor built from a library of CNNs with pre-trained weights derived from the IMAGENET1K_V1 dataset using transfer learning is used to extract the features of the training images³:

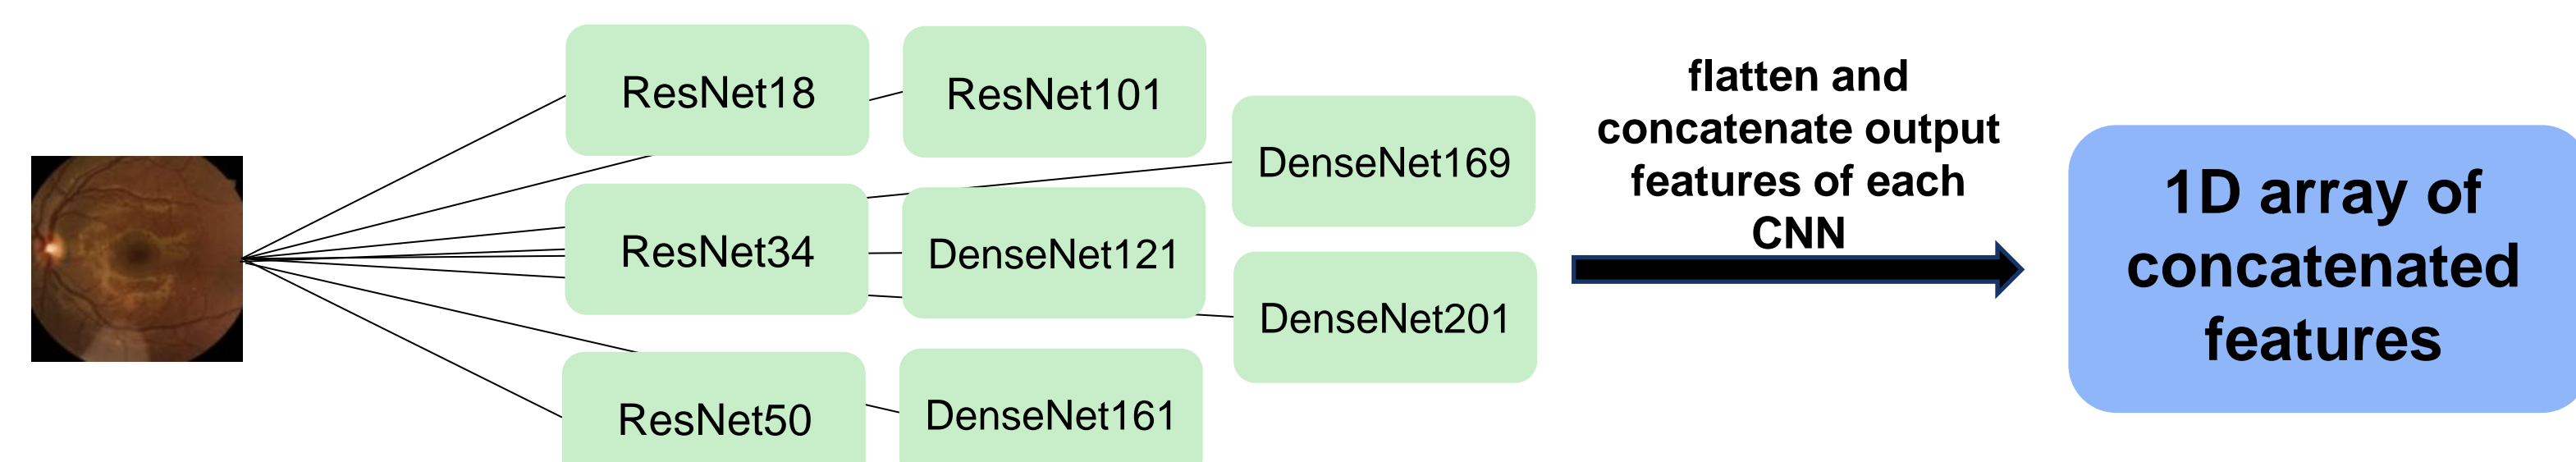


Figure 1 | Architecture of full ensemble feature extractor

Multi-layer Perceptron (MLP) classifier:

The extracted features are passed through an MLP with two hidden layers. The output classifies the input images into one of five DR classes.

Classifier Head Design

Variations in the design of the MLP classifier head has a significant effect on the training dynamics and task performance of the model and plays a large role in the success of training deep neural networks⁴. In this work, we focused on experimentation with variations in the activation function used, the placement of the function within the classifier, and the network's hidden size.

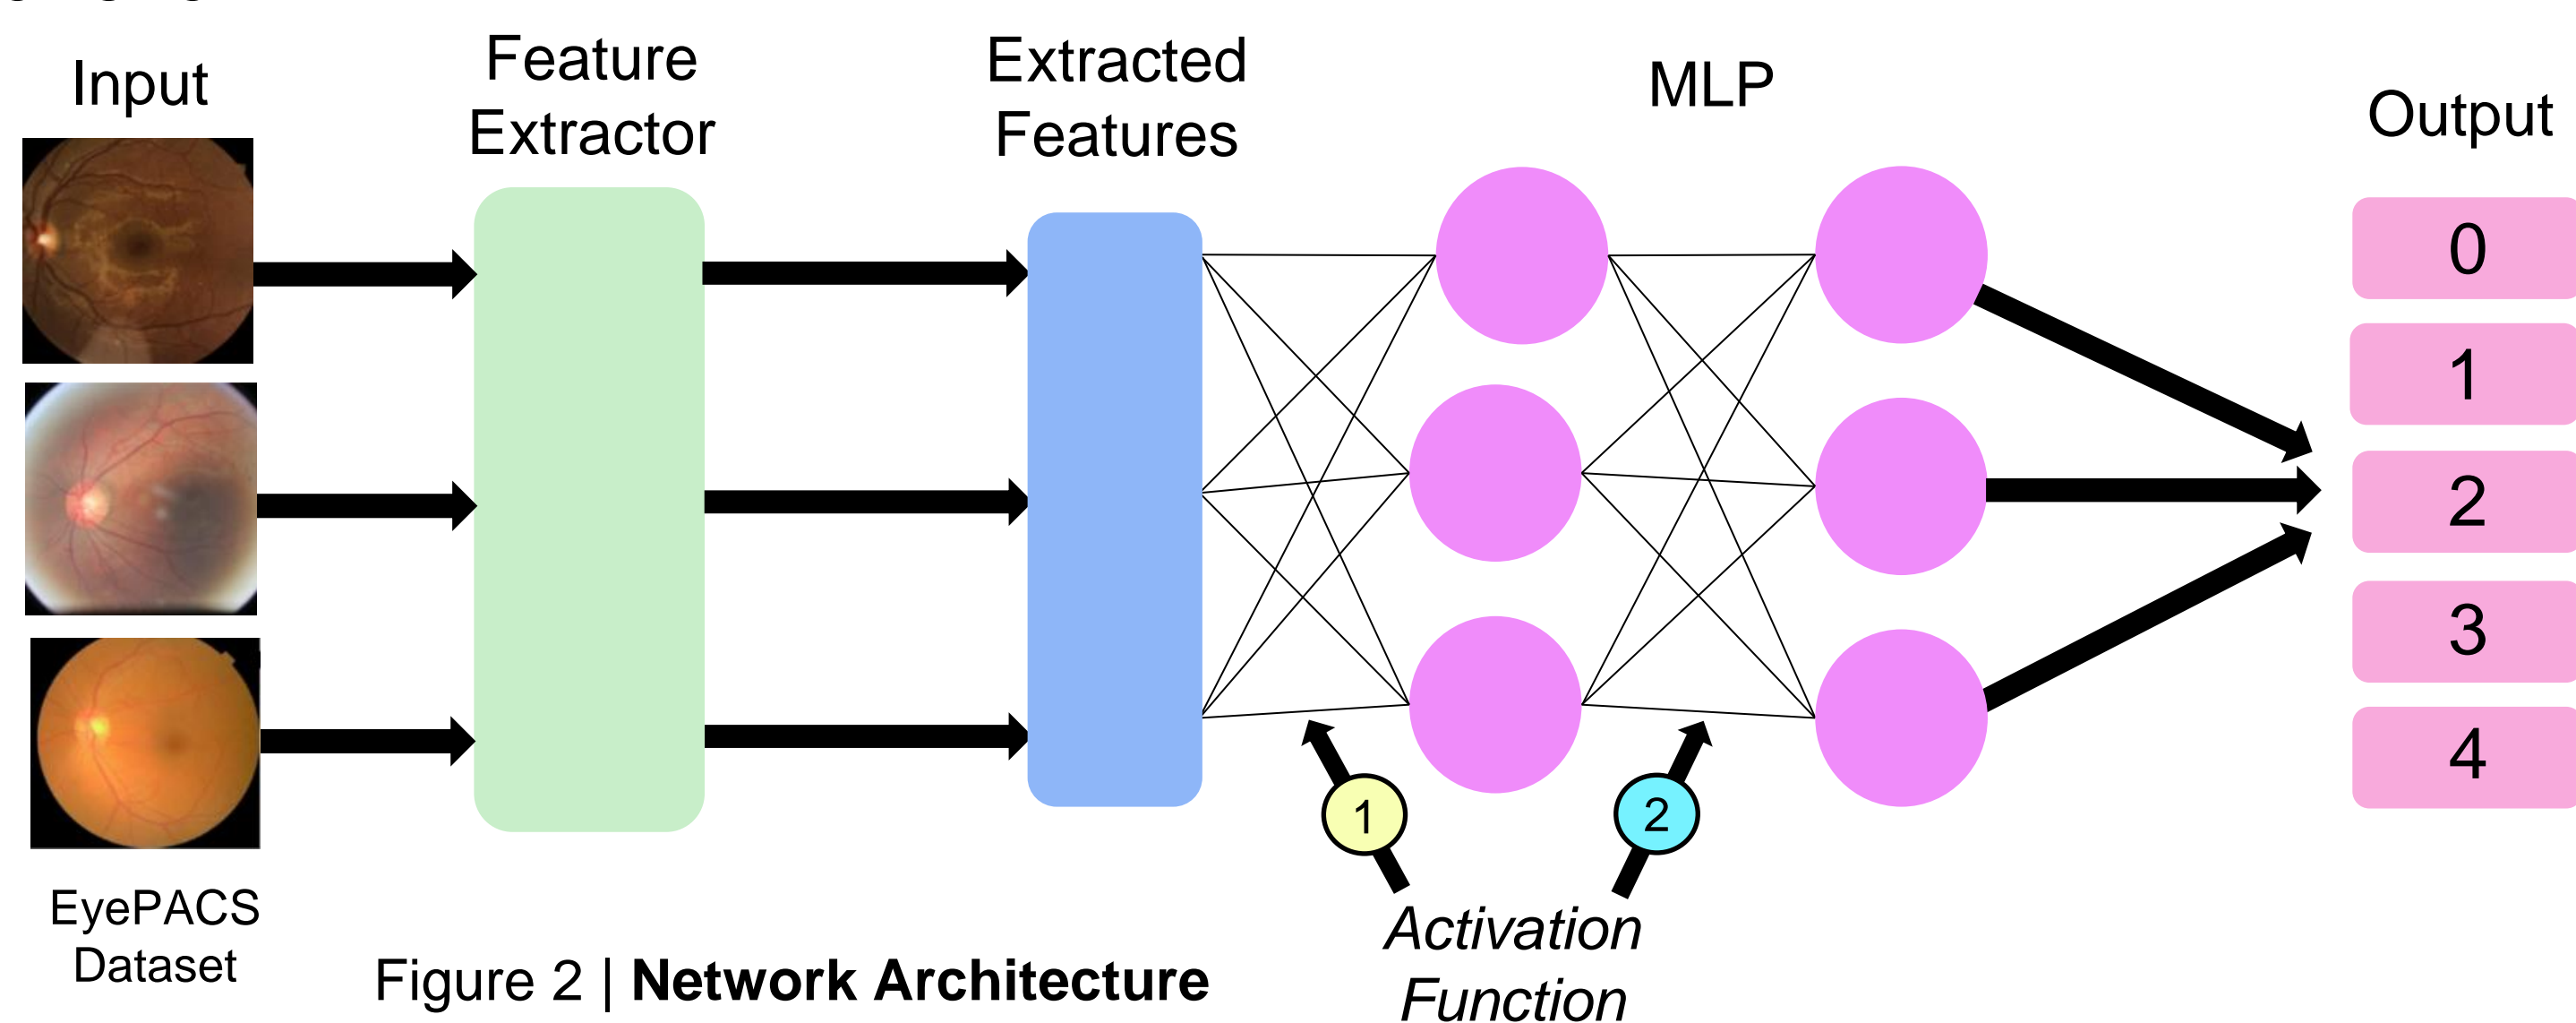
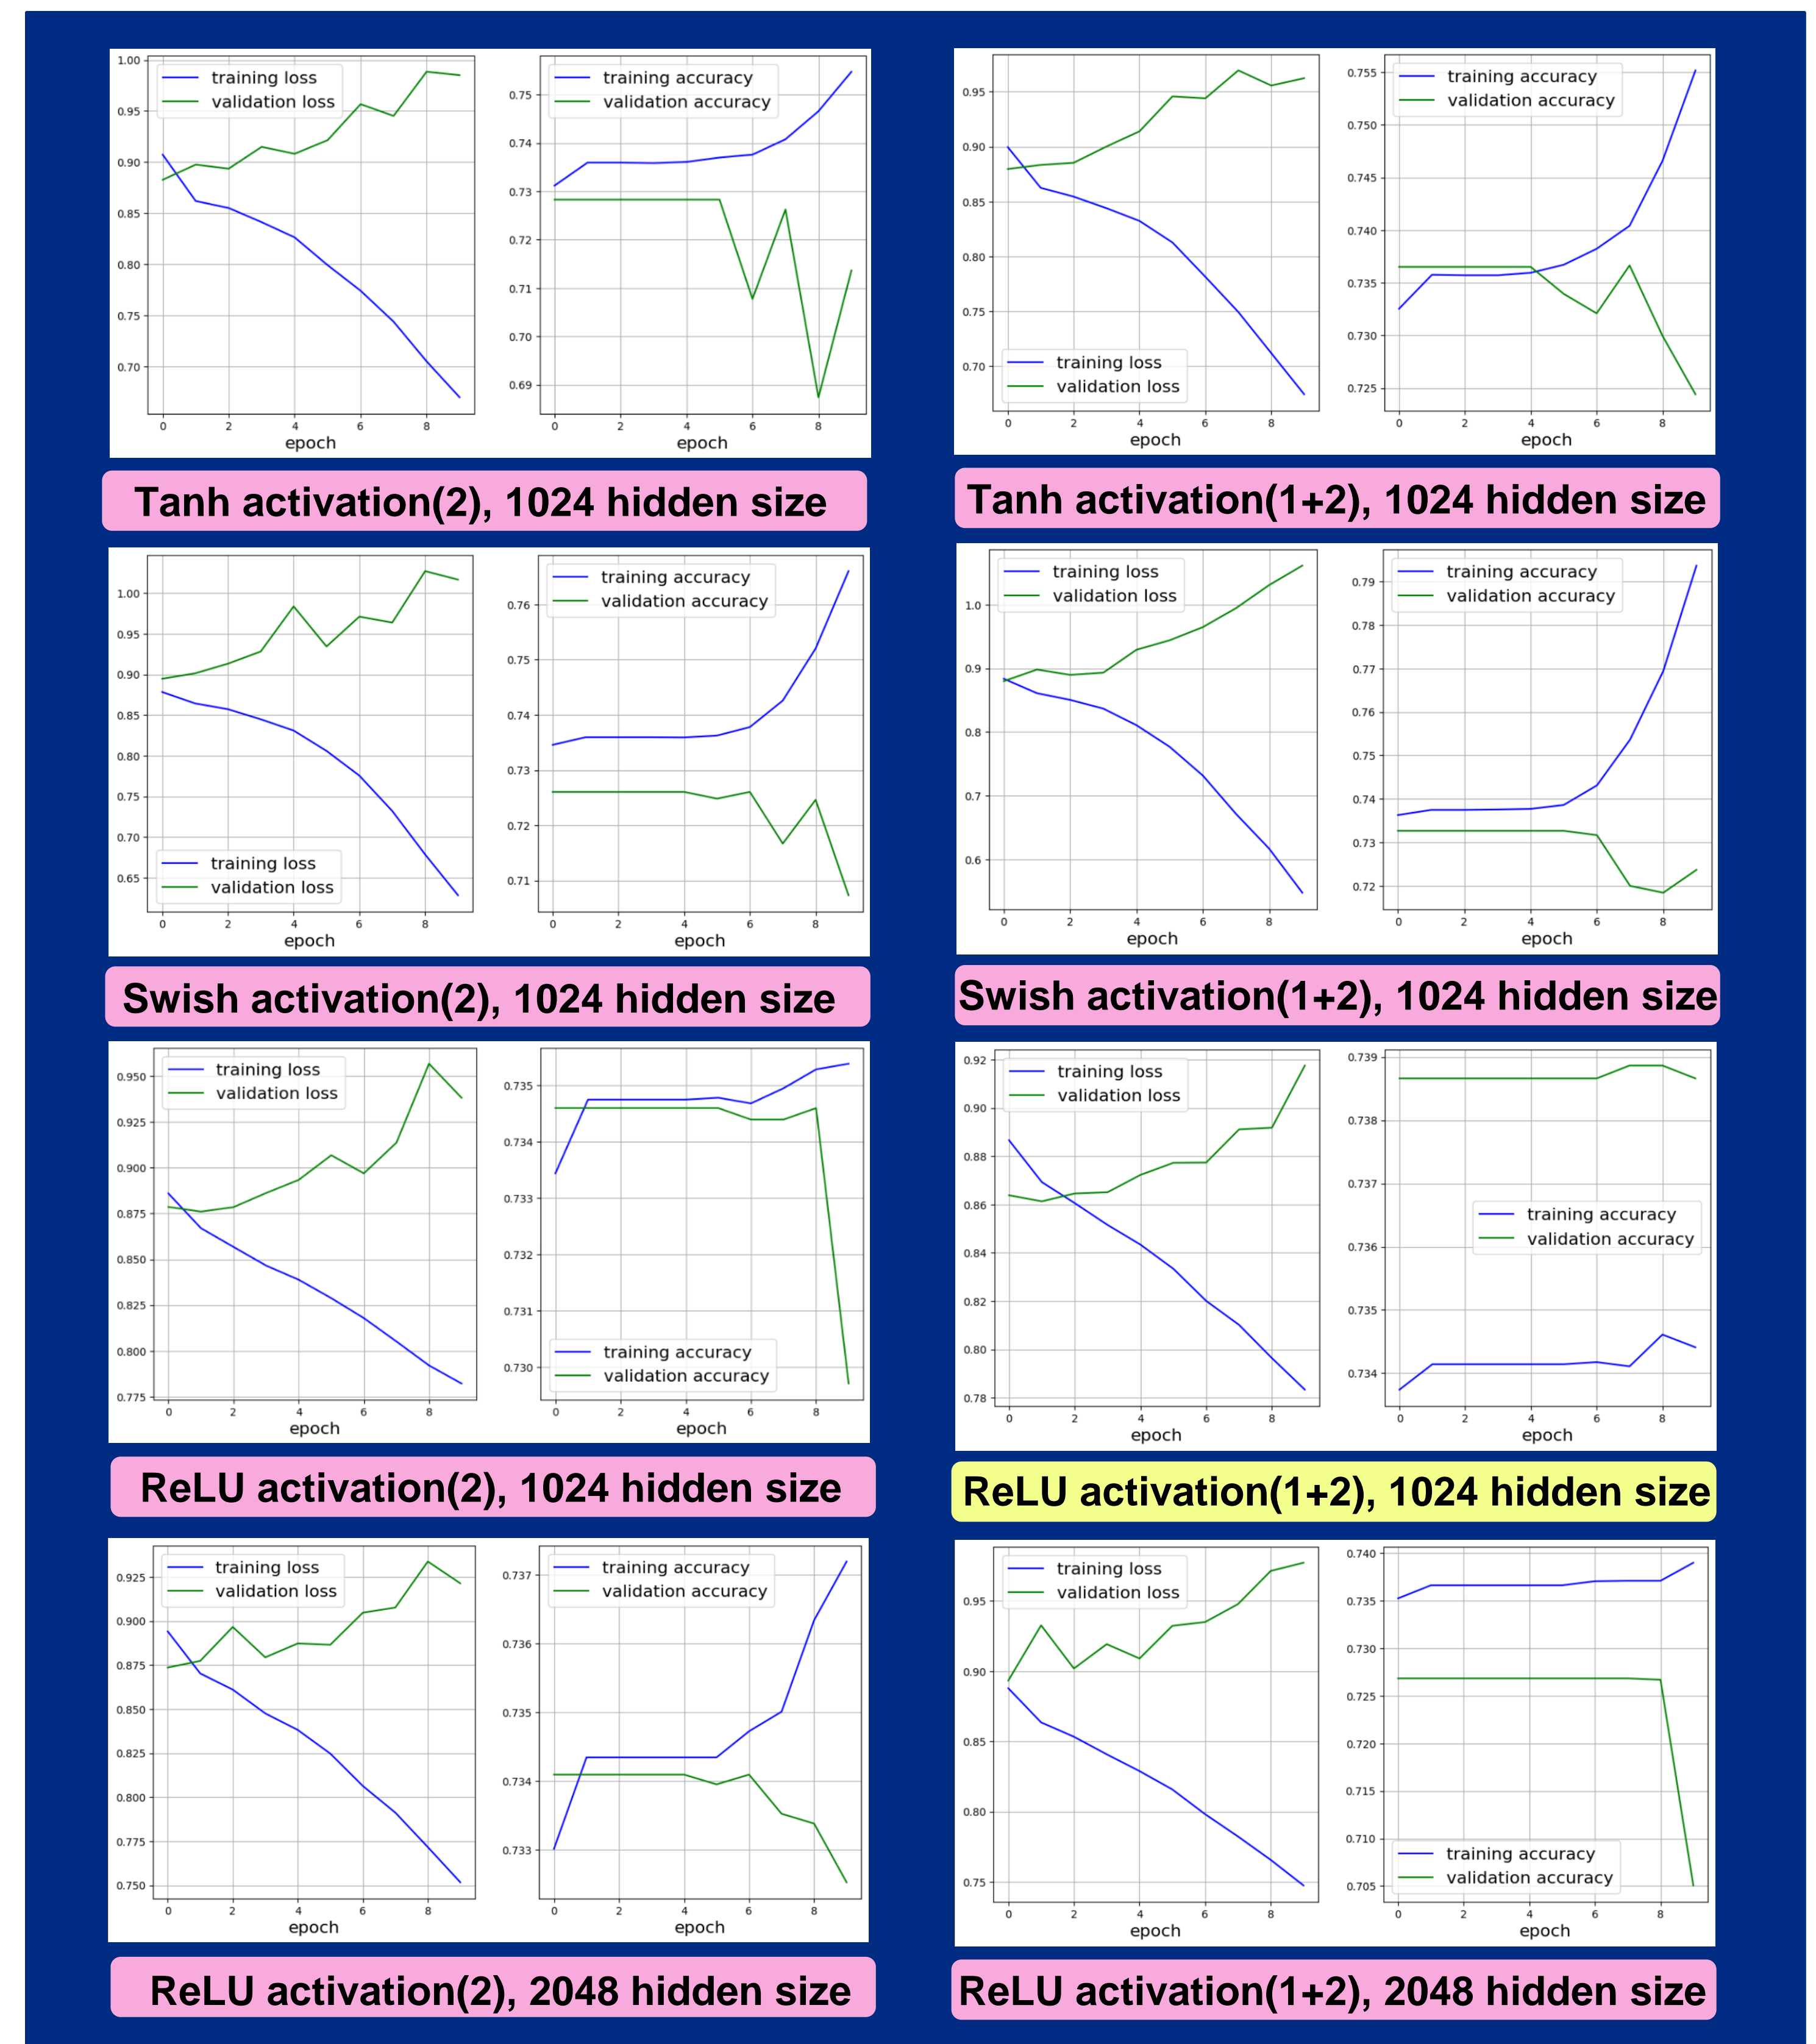


Figure 2 | Network Architecture

Results

In this work, we trained and validated on a subset of the EyePACS dataset using a 60/20/20 train-test-validation split. We experimented with multiple activation functions (TanH, ReLU, Swish), used once or twice in the MLP classifier head (see Figure 2), and with the hidden size of the network.



Best Performance: ReLU Activation(1+2) with 1024 hidden size
Best Epoch (8): 0.7388651616839537 validation accuracy, 0.8911199105250371 validation loss

Figure 3 | Experiments involving different activation functions

Conclusions

We experimented with a variety of MLP classifier head designs to observe which design improves performance. We found that a number of models began to drastically overfit within the earlier epochs. However, several models had less overfitting and produced better results. We believe that with further experimentation and alterations to the design of the MLP classifier head we will be able to improve the diabetic retinopathy classification model.

Future Work

- Implement stratified K-fold cross validation to help detect overfitting
- Apply findings in a few-shot setting to optimize the classification of diabetic retinopathy when there is minimal data available.
- Experiment with the tuning of other aspects within the classifier: L2 regularization, concatenating the features differently, etc.
- Testing the accuracy of the proposed approach on health conditions outside of diabetic retinopathy and eye diseases.
- Investigate what aspects of the image data contribute the most to the model's classification decision.
- Experiment with different combinations of CNNs when creating a full-ensemble of feature extractors.

Acknowledgements

Support for the program REU Site: Scientific Computing for Structure in Big or Complex Datasets is by the NSF grant CNS-1949972, program solicitation NSF 13-582, Research Experiences for Undergraduates (REU). This material is based upon work supported by the National Science Foundation under Grant No. CNS-1949972. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

1. Selvachandran, G., Quek, S.G., Paramesran, R. et al. Developments in the detection of diabetic retinopathy: a state-of-the-art review of computer-aided diagnosis and machine learning methods. *Artif Intell Rev* 56, 915–964 (2023). <https://doi.org/10.1007/s10462-022-10185-6>
2. Diabetic Retinopathy Detection-Kaggle, Oct. 2019, [online] Available: <https://www.kaggle.com/c/diabetic-retinopathy-detection/data>.
3. A. Chowdhury, M. Jiang, S. Chaudhuri and C. Jermaine, "Few-shot Image Classification: Just Use a Library of Pre-trained Feature Extractors and a Simple Classifier," 2021 IEEE/CVF International Conference on Computer Vision (ICCV), Montreal, QC, Canada, 2021, pp. 9425-9434, doi: 10.1109/ICCV48922.2021.00931.
4. Ramachandran, Prajit, Barret Zoph, and Quoc V. Le. "Searching for activation functions." arXiv preprint arXiv:1710.05941 (2017).