Diagnosing diabetic retinopathy from images of the eye fundus

Lucy Wang, Amelie Schäfer {lwang8, amesch}@stanford.edu

https://youtu.be/PzuAppZm9R8

Introduction

- Diabetic retinopathy (DR) is leading cause of blindness in working aged adults
- Need for regular screening without specialist in rural areas
- Use convolutional neural network to classify eye fundus images into 5 disease stages

Dataset

- 3662 labeled 3-channel images from APTOS 2019 Blindness detection challenge on Kaggle [1]
- 5 classes for different disease stages
- Imbalanced class prevalence
- Split into train/dev/test: 70/15/15 preserving class distribution

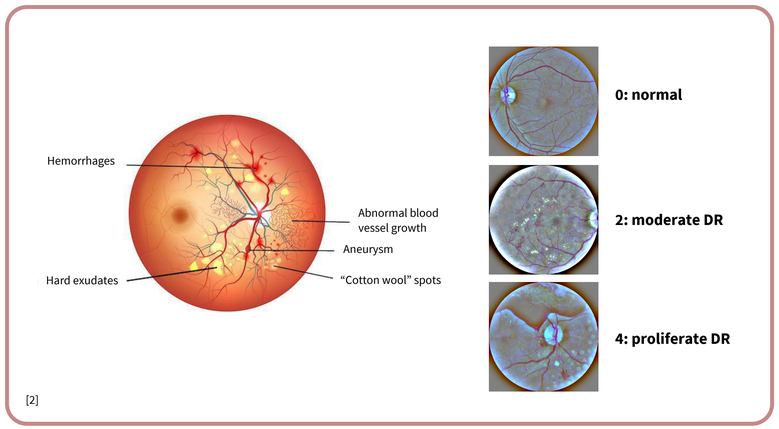

Model

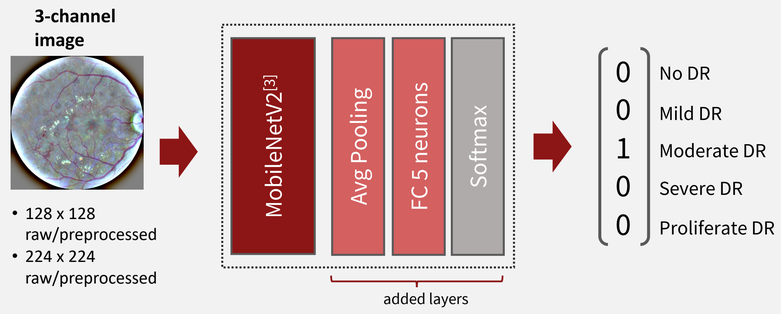

Methods

- Categorical cross entropy loss
- Adam optimizer
- Categorical accuracy
- Quadr. Weighted kappa (QWKP)

Model optimization

- Learning rate
- image size, image preprocessing[4]
- Unfreeze weights
- Weighted loss
- Data augmentation

Results

	Categorical accuracy **Training set**		Categorical accuracy **Development set**	
Initial Raw images 224x224 Model weights frozen	LR 0.001 LR 0.0001 LR 0.00001	0.8246 0.7267 0.6580	LR 0.001 LR 0.0001 LR 0.00001	0.4926 0.6158 0.487
Unfreeze model weights Raw images 224x224 LR: 0.0001	0.9767		Accuracy QWKP	0.6324 0.418
Weighted loss Raw images 224x224 LR: 0.0001	0.9875		Accuracy QWKP	0.5827 0.6833
Data augmentation Rotation, flipping Preprocessed 224x224 LR: 0.0001 Weighted loss	**0.8066**		Accuracy QWKP **Test set** Accuracy QWKP	**0.6930** 0.7761 **0.7847** 0.8123

- Initially high bias and variance
- Unfreezing the pretrained model weights reduces bias
- Data augmentation reduces variance
- Variance still noticeable

Model interpretation

Saliency maps

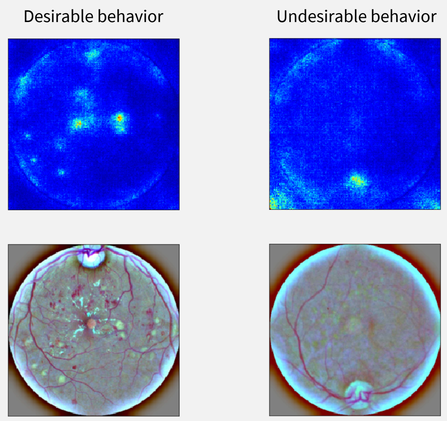

Conclusion

- Challenges: small amount of data, class imbalance
- Strategies: transfer learning, data augmentation, weighted loss
- Performance comparable to lower field of kaggle competitors

Future work

- Cropping background
- L2 regularization
- More data augmentation
- Deeper analysis of saliency maps and class activation maps
- Collect more data

References

[1] https://www.kaggle.com/c/aptos2019-blindness-detection
[2] https://www.eyeops.com/contents/our-services/eye-diseases/diabetic-retinopathy
[3] Sandler, M., Howard, A., Zhu, M., Zhmoginov, A. and Chen, L.C., 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 4510-4520.
[4] https://www.kaggle.com/matheuseduardo/preprocessed-2019-blindness-detection