Segmenting Retinal Blood Vessels with U-Nets

Meeran Ismail (meeran@stanford.edu)

Motivation

Most cases of blindness in American adults occur due to late-stage diabetic retinopathy [3]. Analyzing blood vessels in retinal images is key for early diagnosis of diabetic retinopathy, but is often difficult to discern in blurry retinal images. Automated segmentation can make vessels stand out, and could thus aid less experienced physicians in diagnosing the disease.

Data

Original Data: Training and testing data come from the DRIVE database [2]. There are 20 training examples and 20 testing examples; each example consists of a retinal photo, and an image masks that delineates where the blood vessels are. Each photo and image mask is 565x584 pixels, but are each resized to 512x512 before training and testing the model.

Retinal photo (left), image mask with segmented blood vessels (right)

Augmented Data: To prevent overfitting, we needed more than just 20 pictures. We decided to create two augmented images for each original training picture (40 additional images). We augmented pictures through horizontal flipping, and very slight rotation/zooming.

Augmented image generated by horizontally flipping above photo

U-Net Architecture [5]

Loss Function:
Binary Cross-Entropy

Evaluation Metric:
Pixel Accuracy

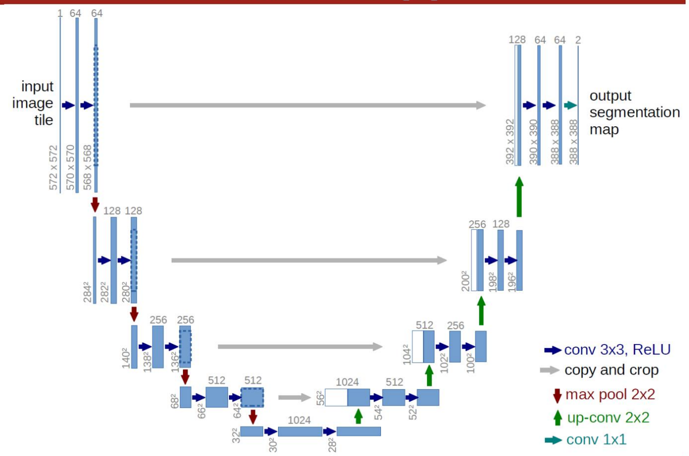

Results

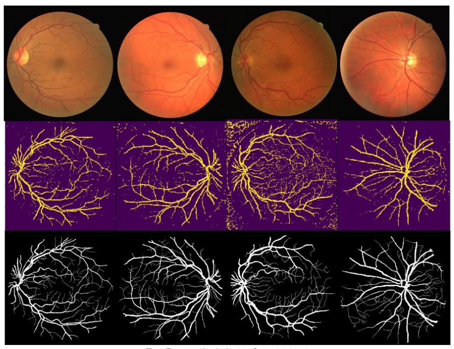

Top Row: retinal photos from test set
Middle Row: model-generated image masks
Bottom Row: ground truth image masks
(each column is one test example)

Model	*Test Accuracy*	*Training Loss*	*Test Loss*
U-Net	**0.9640**	**0.0485**	**0.1438**
Sliding Window ConvNet [4]	0.9013	-	-
SVMs [1]	0.9510	-	-
Adaptive Thresholding [1]	0.9328	-	-
Hierarchal Markovian (Unsupervised) [1]	0.9439	-	-
Random Forest on Gabor Filter Features [1]	0.9464	-	-
Gabor Filter + Fractional Gaussian Derivatives [1]	0.9503	-	-

Previous Approaches

Previous approaches [1] include:

- Sliding window ConvNet
- SVMs
- Adaptive Thresholding
- Unsupervised Hierarchal Markov
- Random Forest on Gabor Filter Features
- Gabor Filter + Fractional Gaussian Derivatives (Unsupervised, see below)

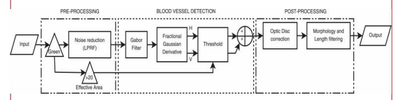

Discussion/Further Work

Judging by raw accuracy, the U-net outperforms all other state-of-the-art models for segmenting blood vessels in retinal photos. However, there are still areas for improvement:

1. Training cross-entropy loss is significantly lower than test loss; increased regularization/dropout/augmentation could help.
2. While false negatives don't appear to be a problem, false positives are, especially in pictures with many faint blood vessels that don't have bright contrast (e.g. 3rd example from the left in results); we can edit the loss function to to more heavily penalize false positives
3. Additionally, there are false positives outside the field of view; perhaps a circular FOV mask can help.
4. We can also try to optimize the DICE coefficient instead

References

[1] Hugo Aguirre-Ramos, Juan Gabriel Avina-Cervantes, Ivan Cruz-Aceves, José Ruiz-Pinales, Sergio Ledesma, "Blood vessel segmentation in retinal fundus images using Gabor filters, fractional derivatives, and Expectation Maximization", *Applied Mathematics and Computation,* 2018, vol. 339, pp. 568-587.
[2] J.J. Stall, M.D. Abramoff, M. Niemeijer, M.A. Viergever, B. van Ginneken, "Ridge based vessel segmentation in color images of the retina", *IEEE Transactions on Medical Imaging,* 2004, vol. 23, pp. 501-209
[3] Klein R, Klein B. "Vision disorders in diabetes." In: National Diabetes Data Group, ed. Diabetes in America. 2nd ed. Bethesda, MD: National Institutes of Health, National Institute
[4] M. Savu, D. Popescu and L. Ichim, "Blood vessel segmentation in eye fundus images," *2017 International Conference on Smart Systems and Technologies (SST)*, Osijek, 2017, pp. 245-249.
[5] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. Unet: Convolutional networks for biomedical image segmentation. CoRR, abs/1505.04597, 2015. URL http:// arxiv.org/abs/1505.04597.