Głębokie sieci konwolucyjne i predykcja strukturalna dla problemów segmentacji naczyń krwionośnych w medycznych obrazowaniach okulistycznych

Paweł Liskowski, Krzysztof Krawiec

Grupa Inteligencji Obliczeniowej
Instytut Informatyki, Politechnika Poznańska

Uczenie maszynowe w robotyce i przetwarzaniu obrazów
Komitet Automatyki i Informatyki PAN, Politechnika Poznańska
18 maja 2018

Diagnosing anomalies in human eye

- Condition of blood vessels in the retina determines the quality of vision.
- Diabetes:
  - 9.3% incidence in the US (2012); over 25% in those aged 65 or more,
  - 7th leading cause of death in the US
  - Diabetic retinopathy affects over a quarter of diabetics
- Medical imaging modalities:
  - Fundus imaging: RGB image of the back of the eye (flash)
  - Fluoresceine angiography (contrast agent required),
  - Optical coherence tomography (OCT).


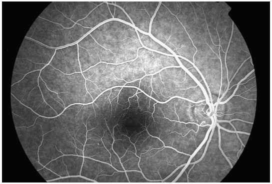
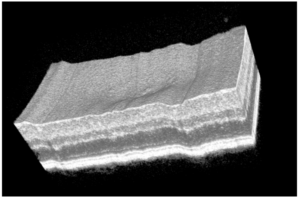

Vessel segmentation problem in fundus imaging

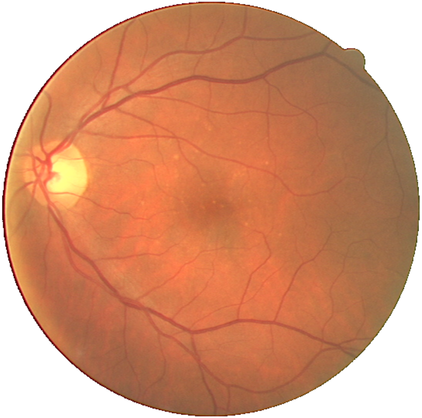

Fundus image

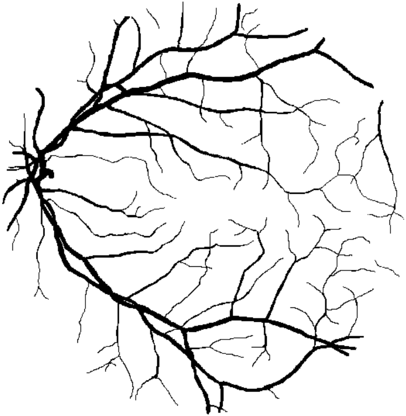

Manual segmentation

Benchmarks datasets for fundus imaging

STARE database
(Hoover et al. 2000)

- TopCon TRV-50 fundus camera
- 20 subjects (images)
  - Includes pathological cases
- No division for training and test sets

DRIVE database
(Staal et al. 2004)

- 3CCD fundus camera,
- 40 subjects
  - 7 with mild diabetic retinopathy
  - 33 clinical norm
- Training set: 20 images.
- Test set: 20 images.

Segmentation task $\Longrightarrow$ Classification task

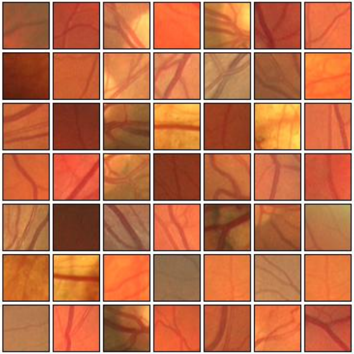

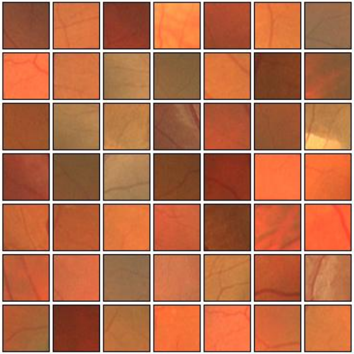

- An example: a 27×27 patch centered at pixel
- The task: Determine the class of the central pixel in a patch
- Binary classification task: vessel vs. non-vessel

Database statistics

	DRIVE		STARE	
Validation scheme	One-off train + test		Leave-one-out	
Image dimensions	565×584		605×700	
	Training	Testing	Training	Testing
Fundus images	20	20	19	1
Total patches	3,857,818	3,993,351	5,326,283	280,330
- positive (vessel)	516,152	538,475	558,261	29,382
- negative	3,341,666	3,454,876	4,768,022	250,948

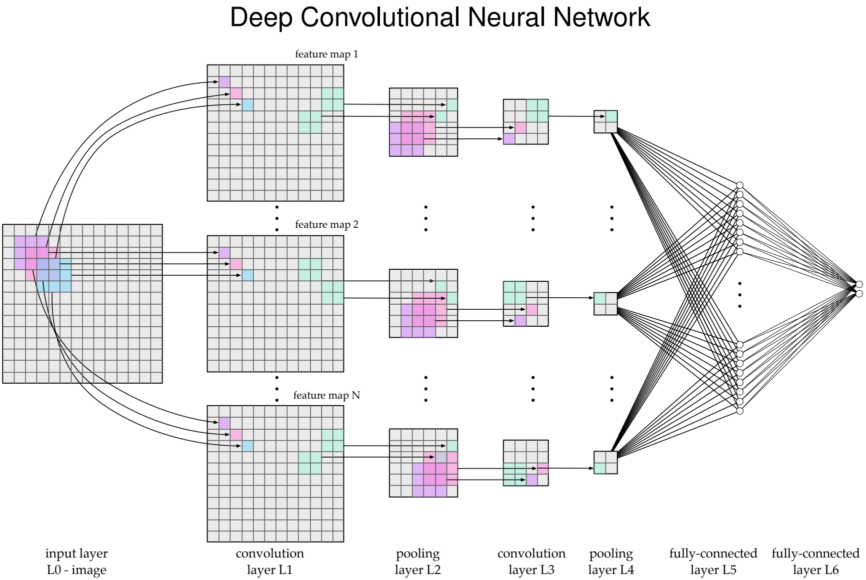
Deep Convolutional Neural Network
feature map 1
feature map 2
feature map N
input layer
L0 - image
convolution
layer L1
pooling
layer L2
convolution
layer L3
pooling
layer L4
fully-connected
layer L5
fully-connected
layer L6

Global contrast normalization

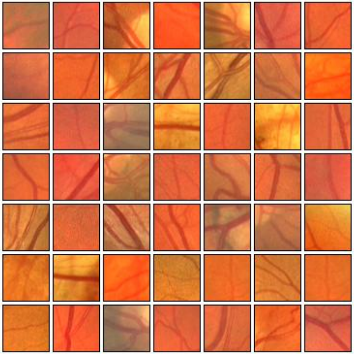

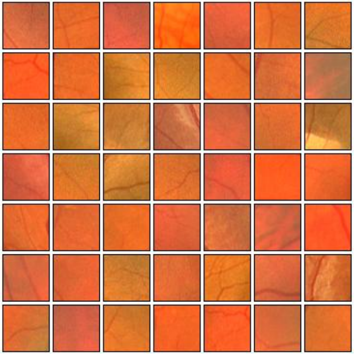

Brightness standardization

Zero-phase Component Analysis (ZCA Whitening)

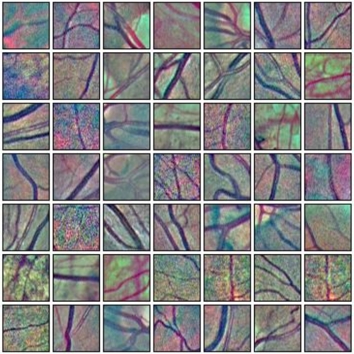

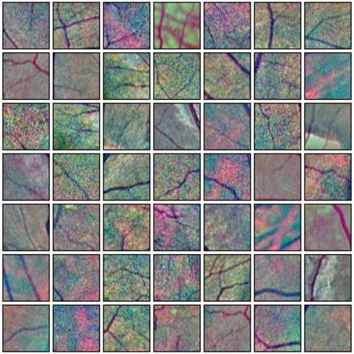

$$W_{ZCA} = U\Lambda^{-1/2}U^T = \Sigma^{-1/2}$$

Augmentations

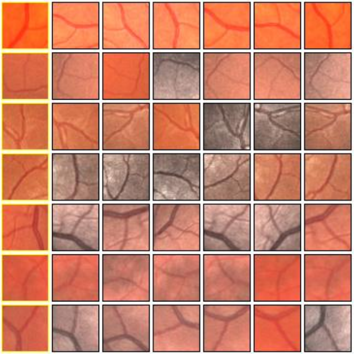

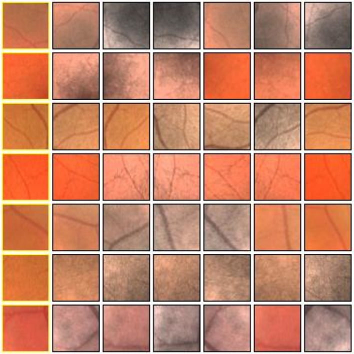

- Scaling by a factor between 0.7 and 1.2,
- Rotation by an angle from $[-90, 90]$,
- Flipping horizontally or vertically,
- Gamma correction
- Overall: 11-fold increase of training set

Parameters and architectures

Training:

- 30,000 iterations (batches) = 19 epochs,
- Rectified Linear Units (ReLUs)
- Learning rate: 10^{-3}, decays 0.1 every 10,000 iterations (batches).
- Momentum term: 0.9,
- Regularization: weight decay (L_2 penalty multiplier 5×10^{-4})
- Dropout ratio: 0.5 (first two FC layers)
- Total number of weights: ~48 millions
- Typical training time: ~10hrs (NVIDIA Titans and Tesla K20c cards).

Testing:

- Querying time: 2.46 ms per pixel, 92 seconds per image (Titan)

Frameworks:

- Caffe, then moved to Theano, plus home-brew glue code in Python

$$\underset{[4\times4\times1\times0]}{conv64} \rightarrow \underset{[3\times3\times1\times1]}{conv64} \rightarrow \underset{[2\times2\times2\times0]}{maxpool} \rightarrow \underset{[3\times3\times1\times1]}{conv128} \rightarrow \underset{[3\times3\times1\times1]}{conv128} \rightarrow \underset{[2\times2\times2\times0]}{maxpool} \rightarrow fc512 \rightarrow fc512 \rightarrow fc2$$

$$\underset{[3\times3\times1\times1]}{conv64} \rightarrow \underset{[3\times3\times1\times1]}{conv64} \rightarrow \underset{[3\times3\times1\times1]}{conv128} \rightarrow \underset{[3\times3\times1\times1]}{conv128} \rightarrow fc512 \rightarrow fc512 \rightarrow fc2$$

Results for STARE

	AUC	Acc	Acc*	Kappa	Sens	Spec
PLAIN	$.9767 \pm .0053$	$.9559 \pm .0071$	$.9551 \pm .0072$	$.7477 \pm .0451$	$.7495 \pm .0721$	$.9788 \pm .0081$
GCN	$.9787 \pm .0049$	$.9571 \pm .0064$	$.9572 \pm .0064$	$.7573 \pm .0394$	$.7620 \pm .0656$	$.9789 \pm .0072$
ZCA	$.9783 \pm .0062$	$.9563 \pm .0064$	$.9562 \pm .0066$	$.7598 \pm .0317$	$.7718 \pm .0490$	$.9783 \pm .0055$
AUGMENT	$.9744 \pm .0048$	$.9527 \pm .0068$	$.9512 \pm .0069$	$.7306 \pm .0431$	$.7376 \pm .0720$	$.9769 \pm .0086$
BALANCED	$.9820 \pm .0045$	$.9309 \pm .0107$	$.9620 \pm .0051$	$.7021 \pm .0305$	$.9307 \pm .0274$	$.9304 \pm .0133$
NO-POOL	$.9785 \pm .0066$	$.9566 \pm .0082$	$.9568 \pm .0081$	$.7622 \pm .0415$	$.7867 \pm .0698$	$.9754 \pm .0099$

(test set, mean and .95 confidence interval)

Best-to-date accuracy and AUC – better than 70+ other methods.

Browse Journals & Magazines > IEEE Transactions on Medical ... > Volume: 35 Issue: 11

Segmenting Retinal Blood Vessels With Deep Neural Networks

Sign In or Purchase
to View Full Text

Related Articles

Extraction of rules from artificial neural networks for nonlinear regression

Time delays and stimulus-dependent pattern formation in periodic environments in...

View All

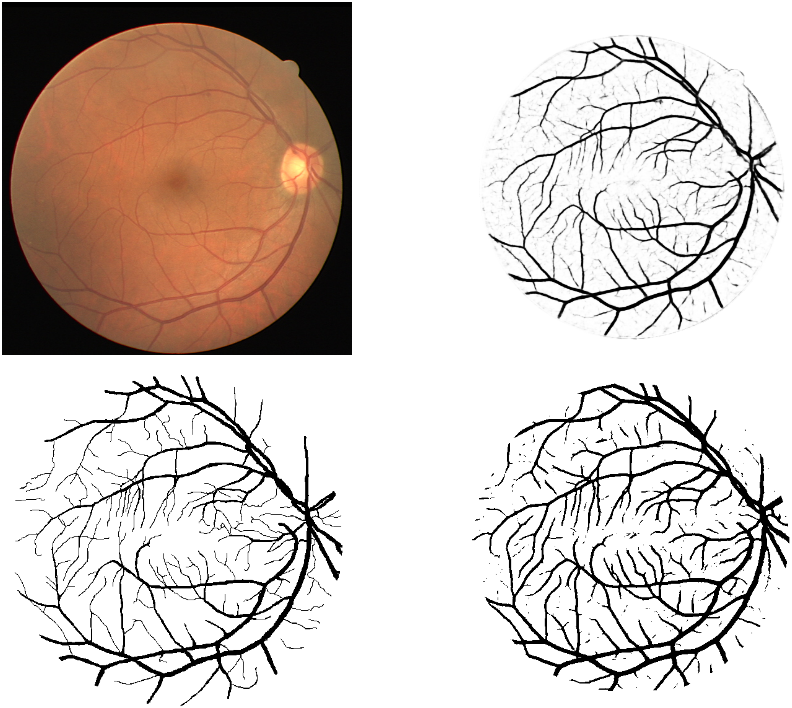

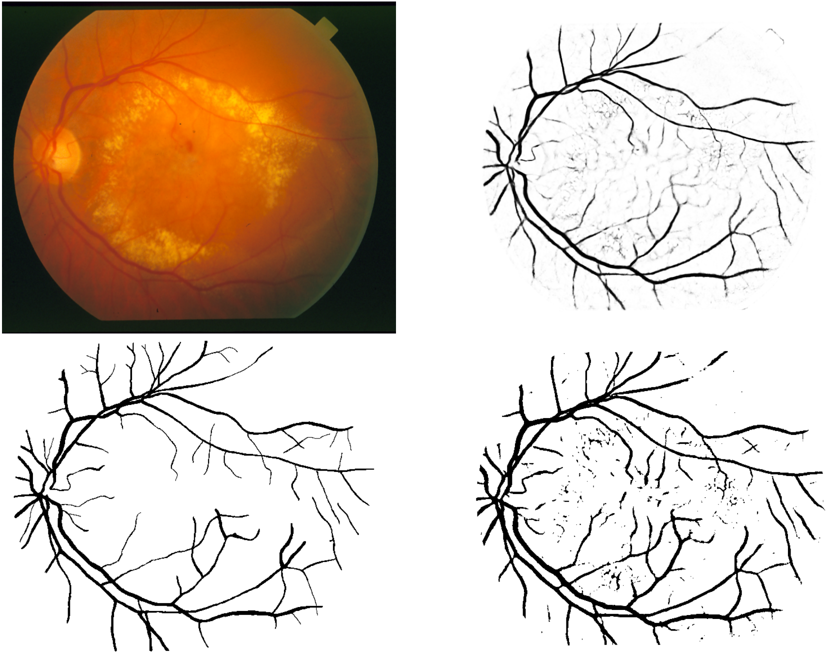

Structured prediction

- Predictions for neighboring pixels are correlated.
- Idea: Classify multiple neighboring pixels (here: 5x5 square).
- Network returns predictions for 25 pixels simultaneously.

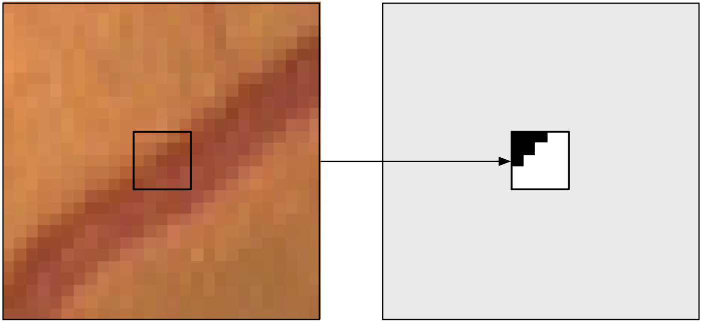

Structured prediction: Results (DRIVE)

Improved Acc and AUC by up to 1%

	s	AUC	Acc	Kappa	Sens	Spec
BALANCED-SP	3	.9787	.9507	.7925	.8460	.9673
NO-POOL-SP	3	.9771	.9519	.7891	.8020	.9757
BALANCED-SP	5	.9788	.9530	.7953	.8149	.9749
NO-POOL-SP	5	.9790	.9535	.7910	.7811	.9807
BALANCED-SP	7	.9729	.9503	.7821	.7996	.9740
NO-POOL-SP	7	.9747	.9518	.7828	.7750	.9795

... and more

- Knowledge transfer
- Central vessel reflex
- Detecting capillaries

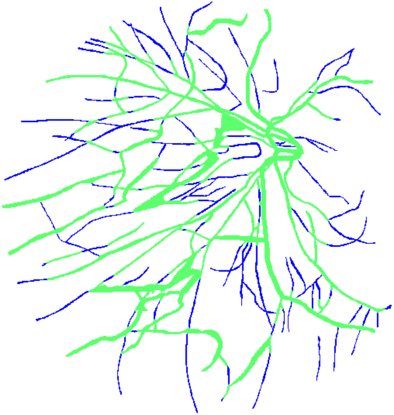

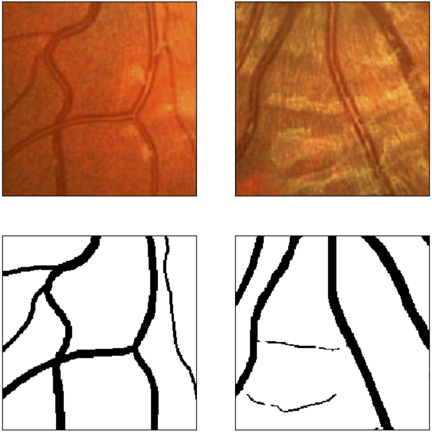

Segmenting 3D OCT images

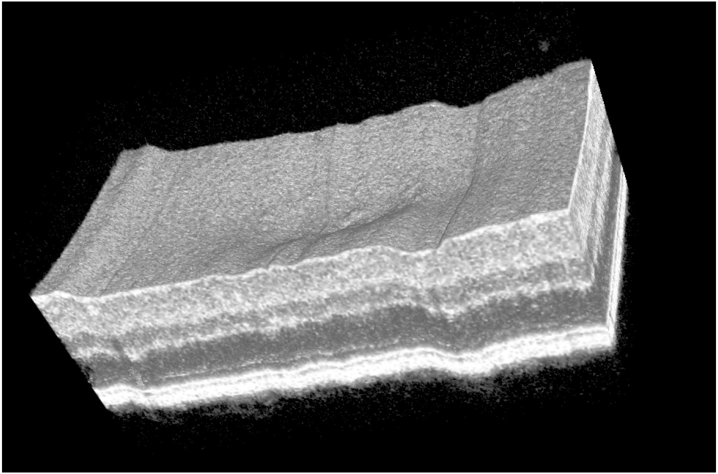

No benchmarks available.
Hard to label.
Lots of data (one image ~ 150MB)

Labeling OCT image

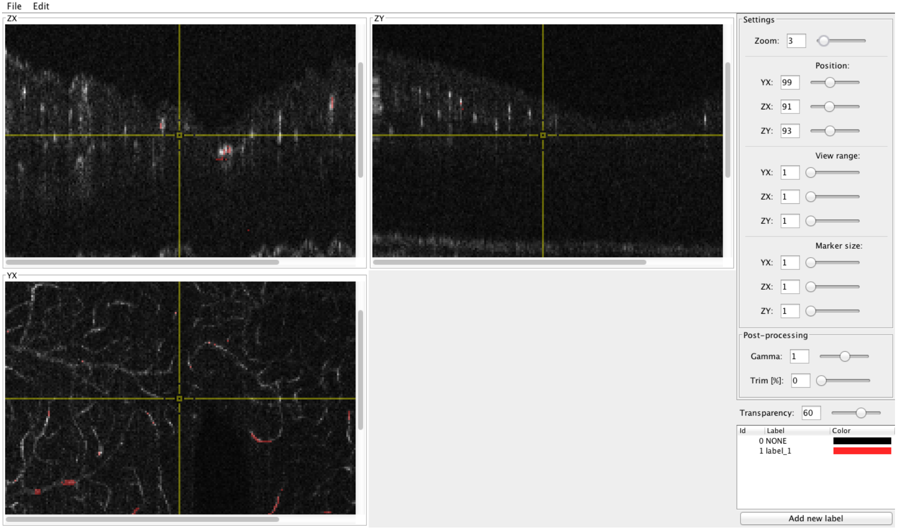

Credits: Bartosz Wieloch

Labeling OCT image

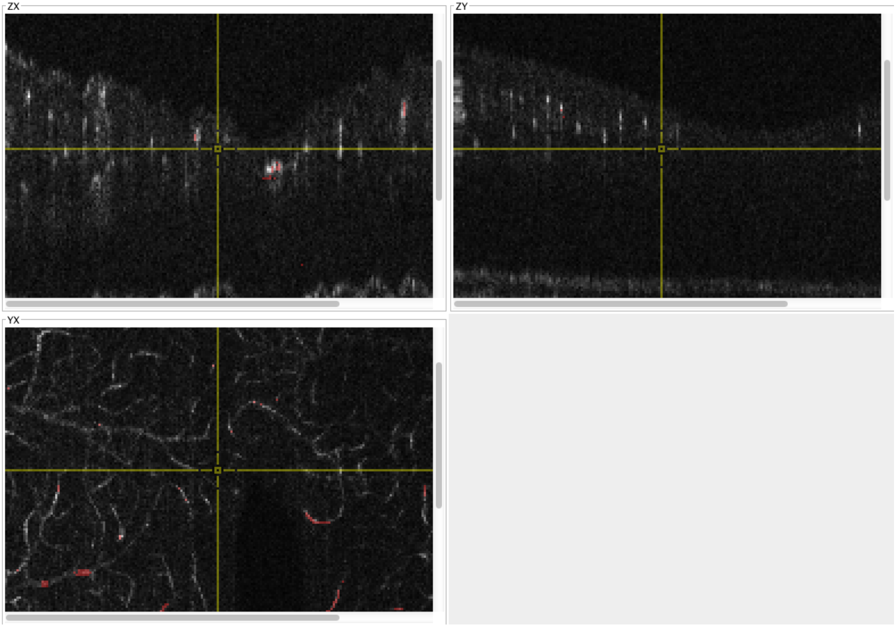

3rd dimension (image ’layers’) interpreted as channels

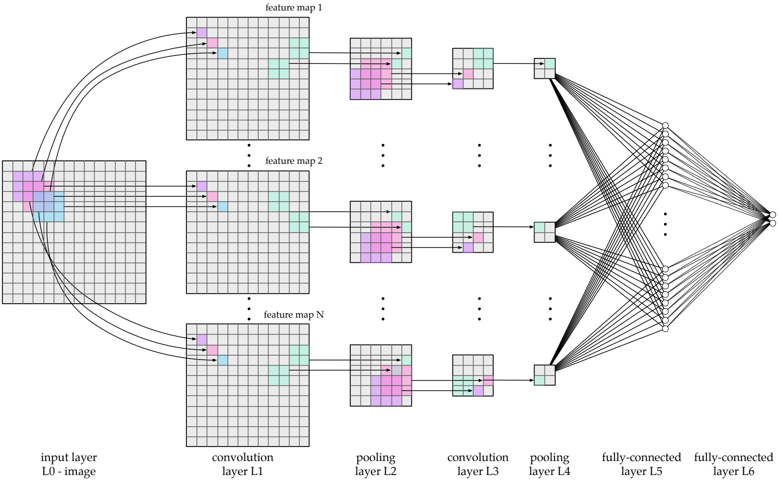

Single example: 27x27x27=19,683 voxels!

Single patch: 27x27 with 27 channels.

Results

Data split	Preprocess.	Model	AUC	Acc	Kappa	Sens	Spec
random	GCN	PLAIN	0.9895	0.9865	0.9831	0.988	0.9851
	SCALING	PLAIN	0.9899	0.988	0.9861	0.9888	0.9873
spatial	GCN	PLAIN	0.9875	0.9722	0.9544	0.9574	0.987
	SCALING	PLAIN	0.9892	0.9773	0.9646	0.9685	0.9861
random	GCN	NOPOOL	0.9898	0.9872	0.9845	0.988	0.9864
	SCALING	NOPOOL	0.9899	0.9875	0.985	0.9883	0.9866
spatial	GCN	NOPOOL	0.9888	0.9753	0.9606	0.9627	0.9879
	SCALING	NOPOOL	0.9882	0.9521	0.9142	0.9154	0.9888

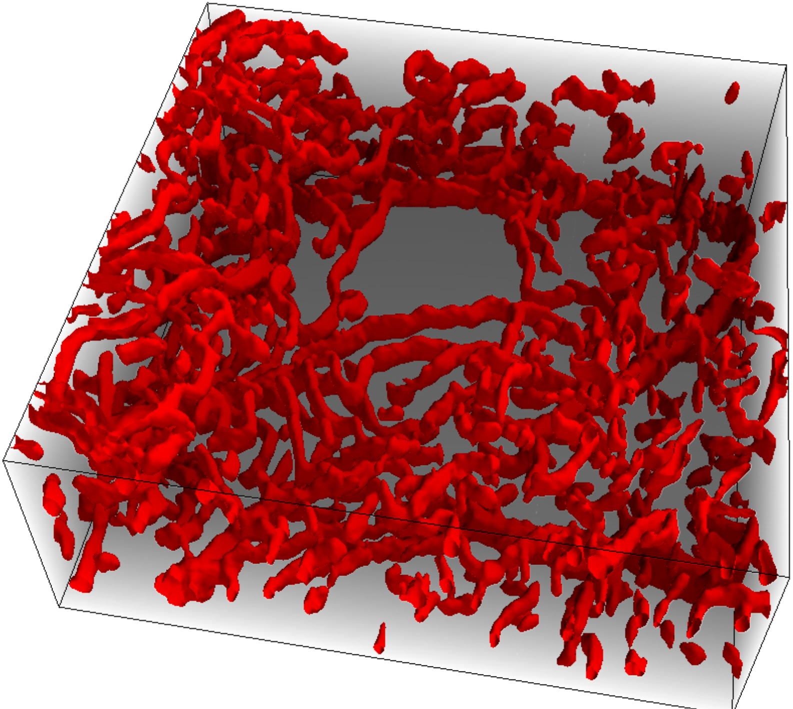

Summary

- Superb accuracy on challenging segmentation tasks.
  - No explicit domain knowledge required.
  - Versatile.
- Challenges: training time, querying time.

Summary

- Superb accuracy on challenging segmentation tasks.
  - No explicit domain knowledge required.
  - Versatile.
- Challenges: training time, querying time.

Credits: Bartosz Wieloch

National Centre for Research and Development:

Development of Interferometric Imaging Methods for Investigation of Dynamics of Biological Systems (2013–2015), NCBiR PBS1/A9/20/2013, 5,926,260PLN

P. Liskowski, K. Krawiec, Segmenting Retinal Blood Vessels with Deep Neural Networks, *IEEE Transactions on Medical Imaging*, vol. 35(11), 2016, pp. 2369-2380. M. Szkulmowski, P.

Liskowski, B. Wieloch, K. Krawiec, and B. Sikorski. Convolutional Neural Networks for Artifact Free OCT Retinal Angiography, The Annual Meeting of the Association for Research in Vision and Ophthalmology, 2017.