Mitosis Detection in Breast Cancer Histology Images with Multi Column Deep Neural Networks

Dan C. Cireşan and Alessandro Giusti

IDSIA, Lugano, Switzerland
dan.ciresan@gmail.com

IDSIA

DNN for Visual Pattern Recognition

- One of the first to have a DNN implemented on GPU (CUDA), 2009
- We applied DNN on a plethora of pattern recognition tasks

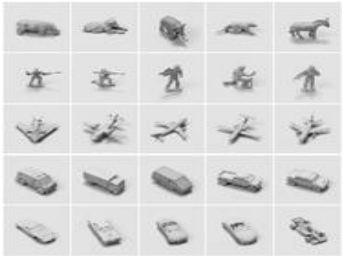

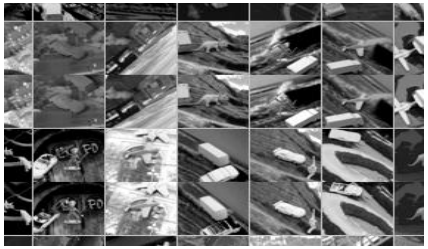

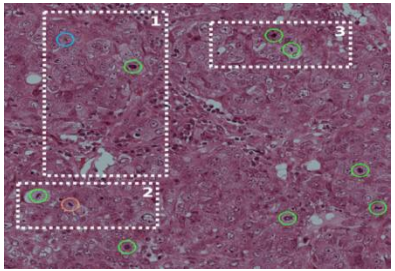

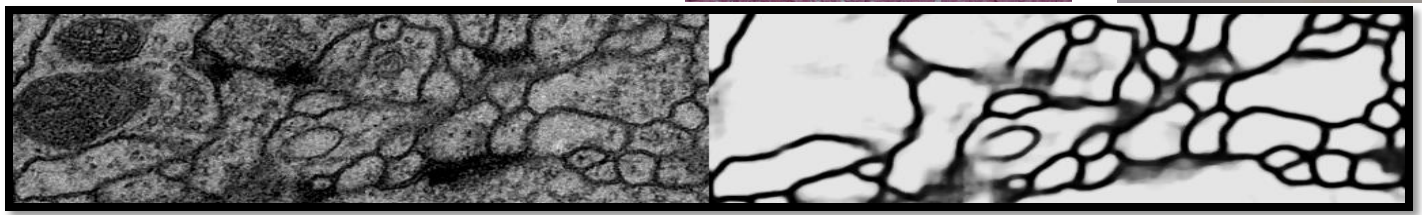

Why mitosis detection?

- Mitosis detection is a challenging visual pattern recognition problem
- No histology or medicine background
- **ICPR2012 & MICCAI2013 competitions:**
  - 2012 ICPR Competition: 50 images, 300 mitosis; 17 teams
  - 2013 MICCAI Competition: ~600 images, 1157 mitosis; 14 teams

Deep, Convolutional Neural Network

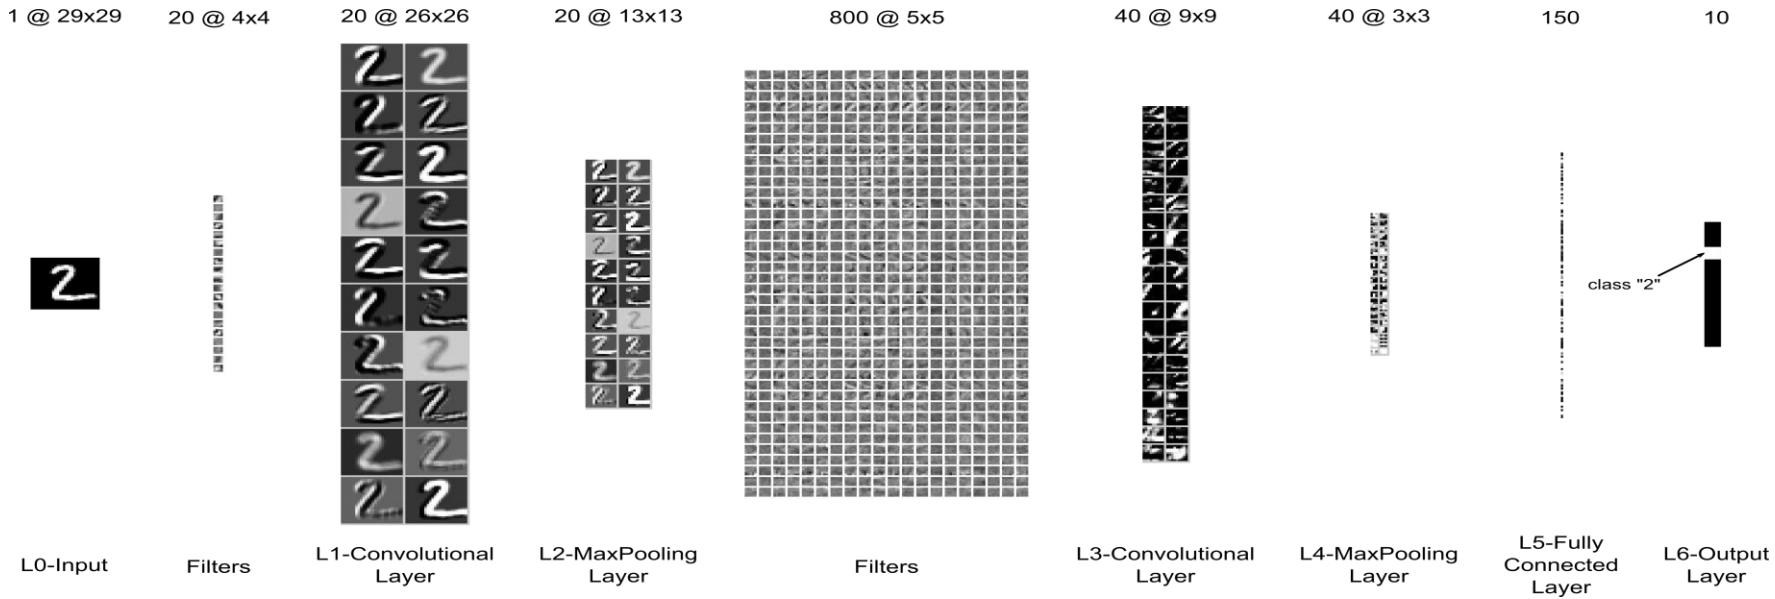

http://ipal.cnrs.fr/ICPR2012/

MITOSIS DETECTION IN BREAST CANCER HISTOLOGICAL IMAGES

IPAL UMI CNRS - TRIBVN - Pitié-Salpêtrière Hospital - The Ohio State University

Home Organisers Important Dates Description Evaluation Metrics Dataset Results Registration

Mitosis Detection in Breast Cancer Histological Images (MITOS dataset)

An ICPR 2012 Contest

We propose a contest of mitosis detection in images of H&E stained slides of breast cancer. Mitotic count is an important parameter for the prognosis of breast cancer. However, mitosis detection is a challenging problem and has not been addressed well in the literature. Indeed, mitosis detection is very challenging since mitosis appear in image as small objects with a large variety of shapes, and they can thus be confused with some other objects or artefacts present in the image.

We add a further dimension to the contest by using two different slide scanners having different resolution to produce RGB images and a multi-spectral microscope producing images in 10 different spectral bands and 17 layers Z-stack.

USER LOGIN

Username *

Password *

Request new password

Log in

D. Ciresan et al. - Mitosis Detection in Breast Cancer Histology Images using Deep Neural Networks, MICCAI 2013

Data Description

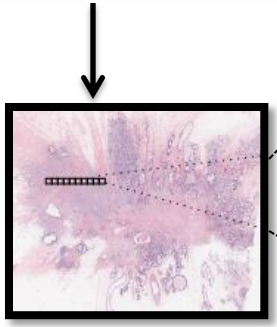

2048x2048 px (0.5 x 0.5 mm)

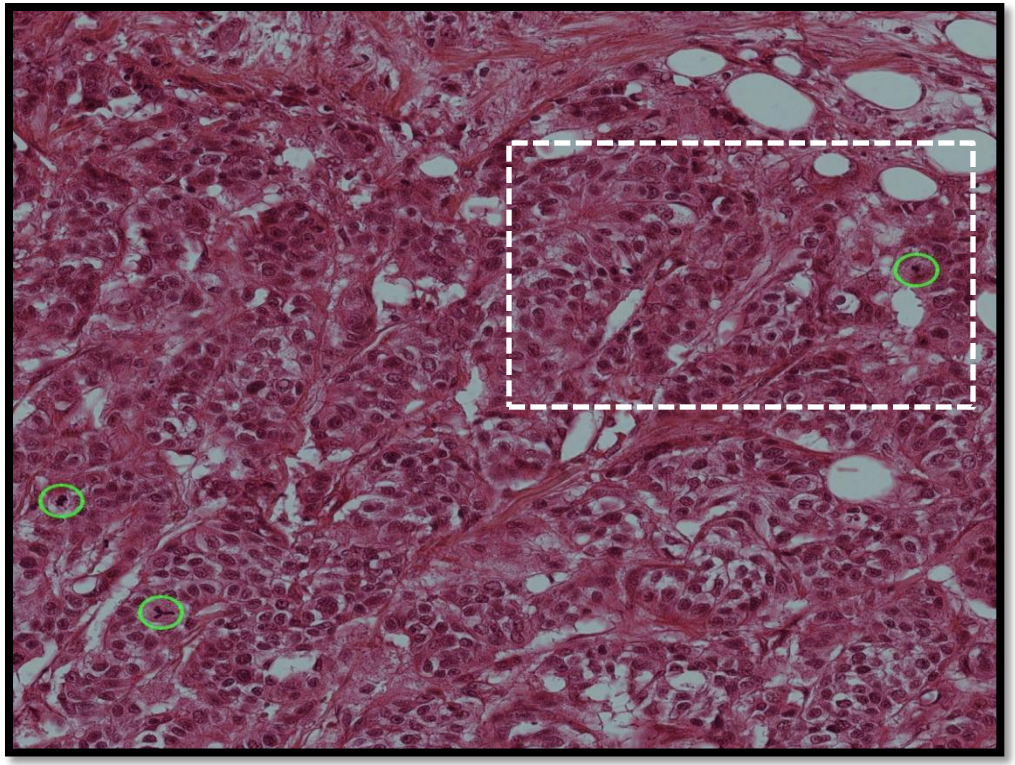

Method

- We use a **powerful pixel classifier** (a Deep Convolutional Neural Network) to detect pixels close to mitosis centroids
- Input: **raw pixel values** in a window (no features, no preprocessing)
- Output: **probability** of central pixel being close to a mitosis centroid

Network Architecture

Layer	Type	Maps and neurons	Filter size	Weights	Connections

Training samples & time

66K positive training samples
(all pixels closer than 10 px to a mitosis)
2M negative training samples

ICPR 2012

→**5 months** training time for up to 7 epochs on a **CPU**

Or up to **3 days** on a **GPU**

Approach Overview

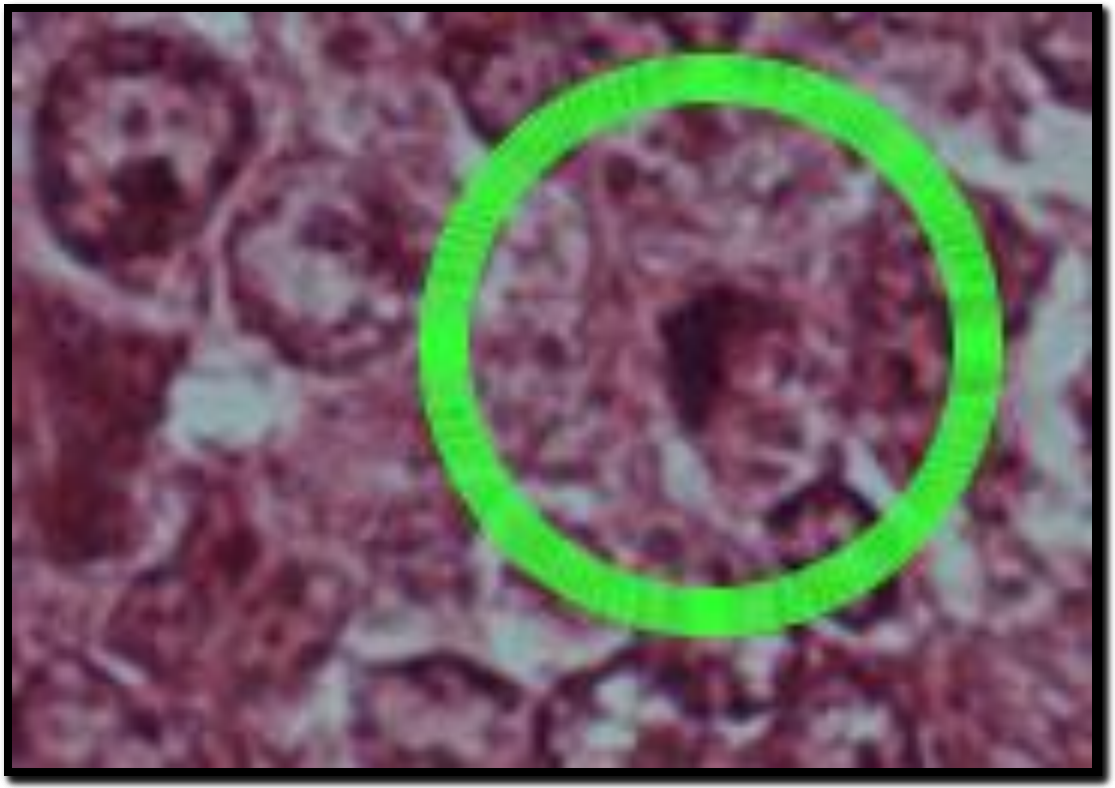

Data and nets

- Training set (263 images with ground truth, coming from 12 patients)
  - We split the training set in two sets T1 (174 images) and T2 (89 images)
- Initially we trained nets on T1 and validated on T2
- Then we trained nets on T1+T2 and applied them to T3 (our submissions)
- Evaluation set (295 images without ground truth, coming from other 11 patients)
  - Used exclusively for testing
  - Denoted as T3 (ground truth known only by the organizers)

Results for net n10

- Trained on T1. Results on the validation set (T2) with 8 variations
- Ground truth is used to decide on which threshold to use when training on T1+T2
- T2: (max F1 ~0.64, F1 at 0.4 ~0.6) T3: F1 at 0.4 0.505
- We either overfitted on T2, or T2 and T3 are quite different (or both)

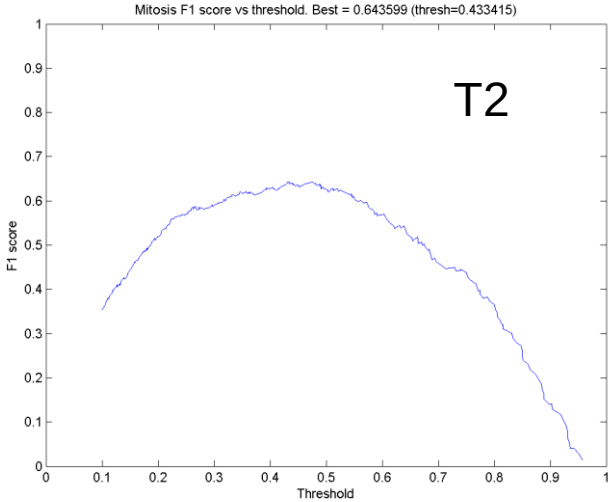

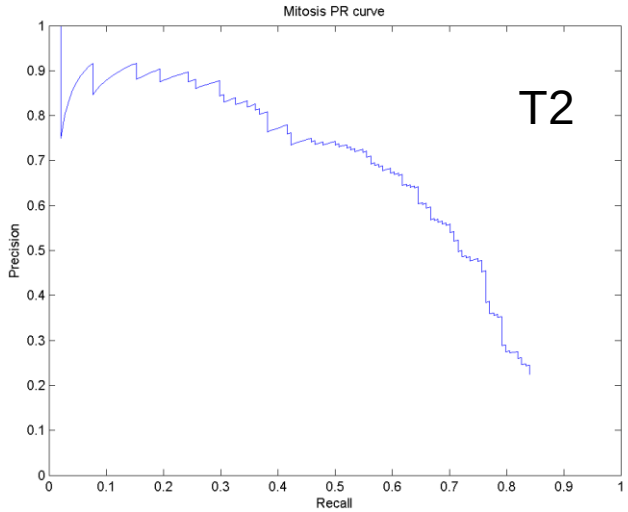

Our submissions

- n10e06 + n30e05 + n31e02, 8 variations, T1+T2
  - t=0.45 -> F1-score = **0.593**
  - t=0.35 -> F1-score = **0.460**
  - t=0.5 -> F1-score = **0.611**
- n10e06, 8 variations, T1, t=0.4 -> F1-score = **0.505**

Results overview on the evaluation dataset

Green: True Positives
Red: False Positives
Cyan: False Negatives

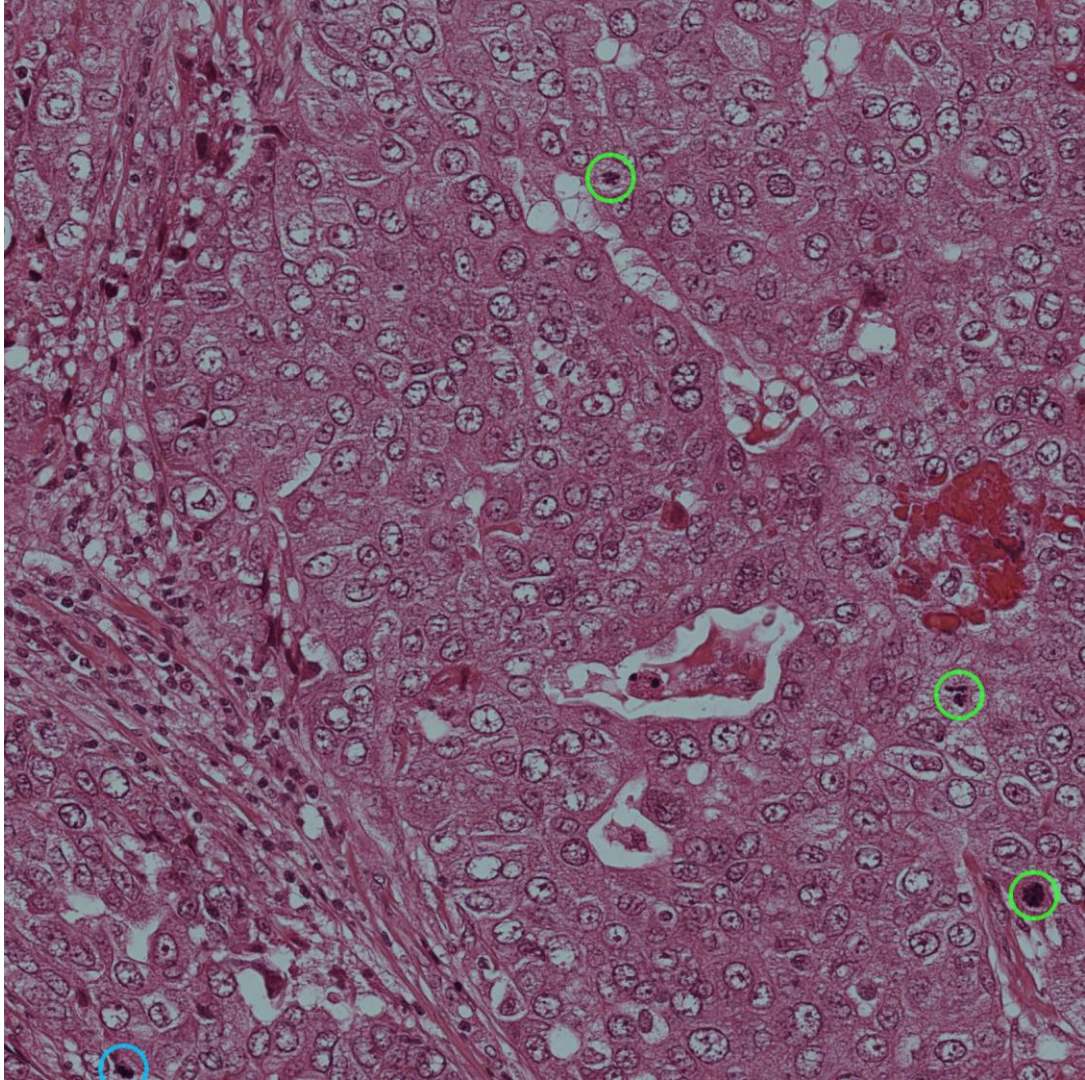

n10 on validation data (T2)

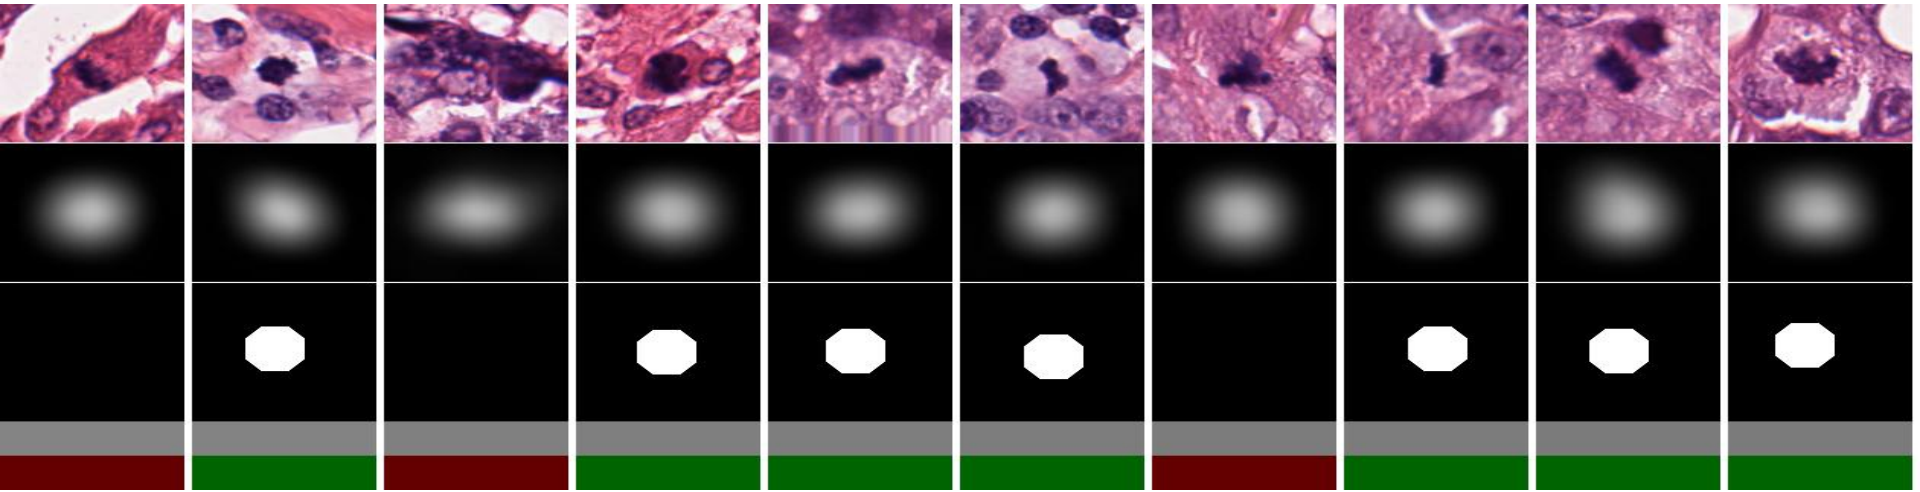

Detection results

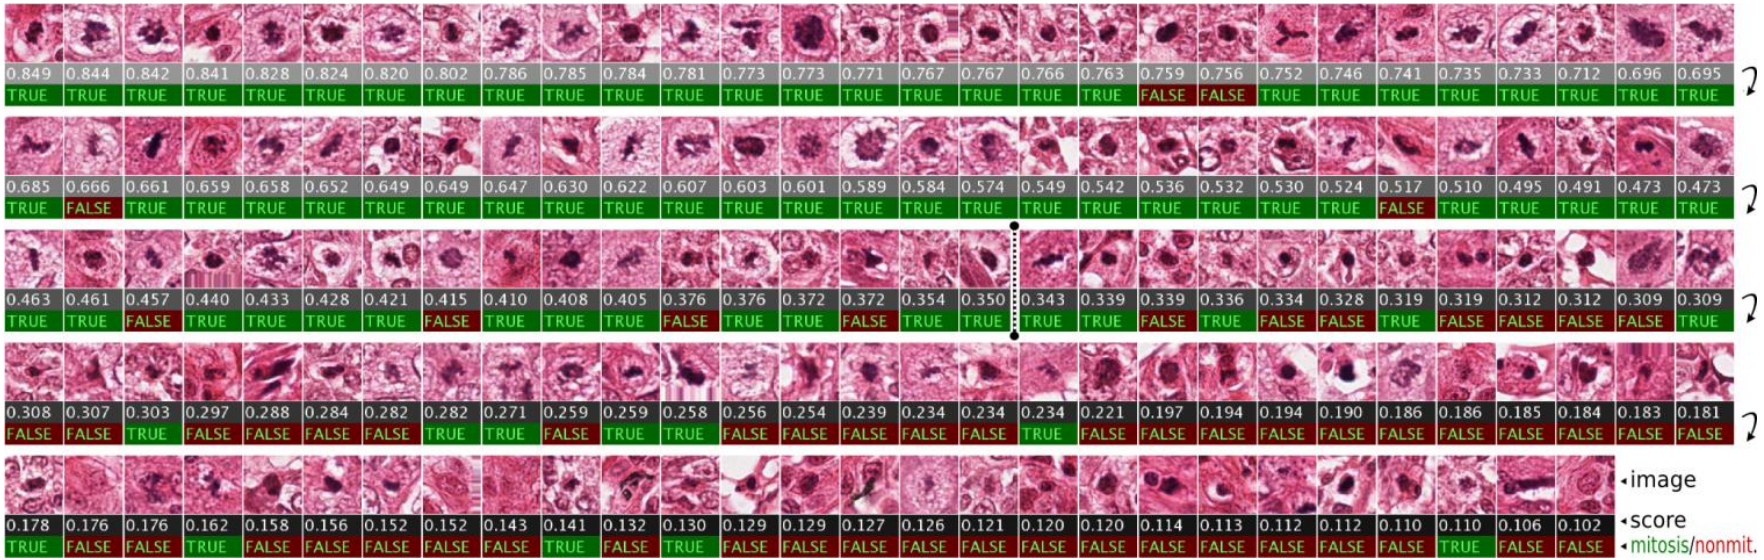

method	precision	recall	F_1 score
DNN	0.88	0.70	**0.782**
DNNf12	0.85	0.68	0.758
DNNf1	0.78	0.72	0.751
IPAL [7]	0.69	0.74	0.718
SUTECH	0.70	0.72	0.709
NEC [12]	0.74	0.59	0.659
UTRECHT [20]	0.51	0.68	0.583
WARWICK [10]	0.46	0.57	0.513

method	precision	recall	F_1 score
NUS	0.63	0.40	0.490
ISIK [19]	0.28	0.68	0.397
ETH-HEILDERBERG [18]	0.14	0.80	0.374
OKAN-IRISA-LIAMA	0.78	0.22	0.343
IITG	0.17	0.46	0.255
DREXEL	0.14	0.21	0.172
BII	0.10	0.32	0.156
QATAR	0.00	0.94	0.005

D. Ciresan et al. - Mitosis Detection in Breast Cancer Histology Images using Deep Neural Networks, MICCAI 2013

Quantitative Results

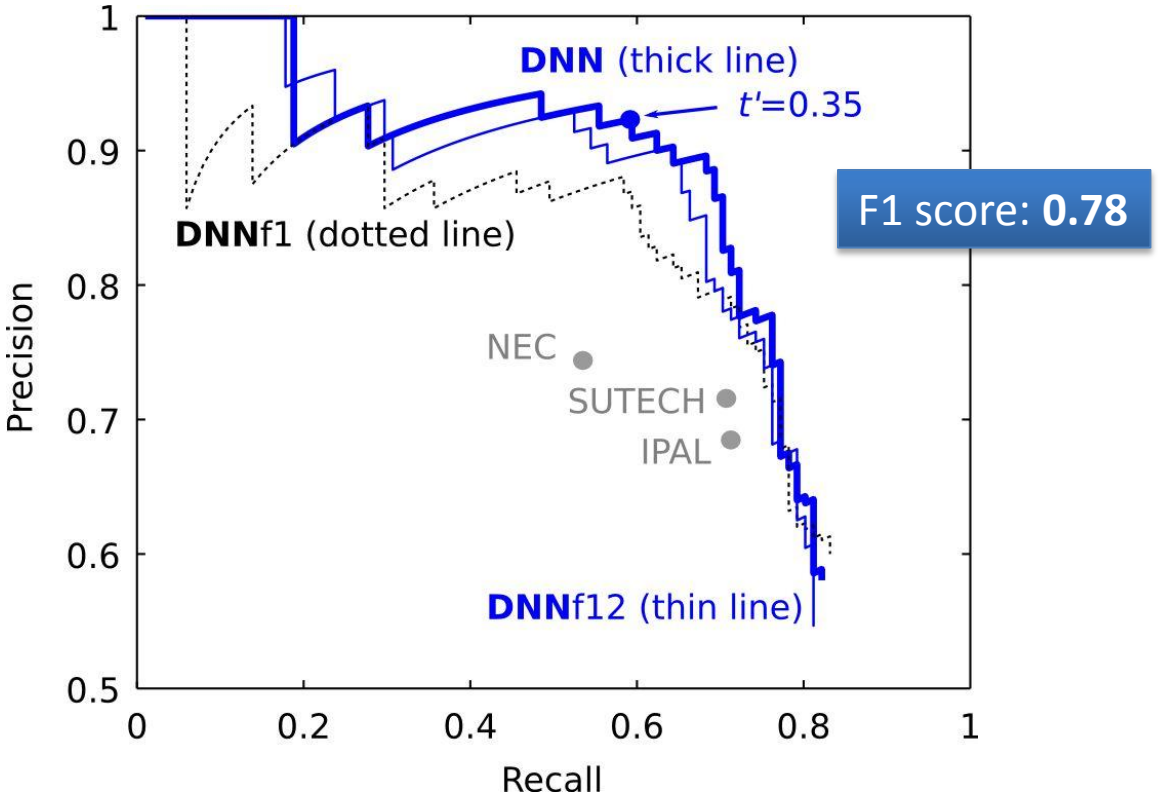

F1 score: **0.78**

Assessment of Mitosis Detection Algorithms 2013 - MICCAI Grand Challenge

Assessment of Mitosis Detection Algorithms 2013

AMIDA13 | MICCAI Grand Challenge

http://amida13.isi.uu.nl/

Home | Background | Dataset | Evaluation | Rules | Results | Tools | Contact

Share 5 Tweet 7 +1 5 Recommend 12

- more training data
  - 2012 ICPR Competition
    - 50 images, 300 mitosis, 17 teams
  - 2013 MICCAI Competition
    - ~600 images, 1157 mitosis, 14 teams
- test data is more difficult

User login

Username *

Password *

- Create new account
- Request new password

Log in

Updates

Results

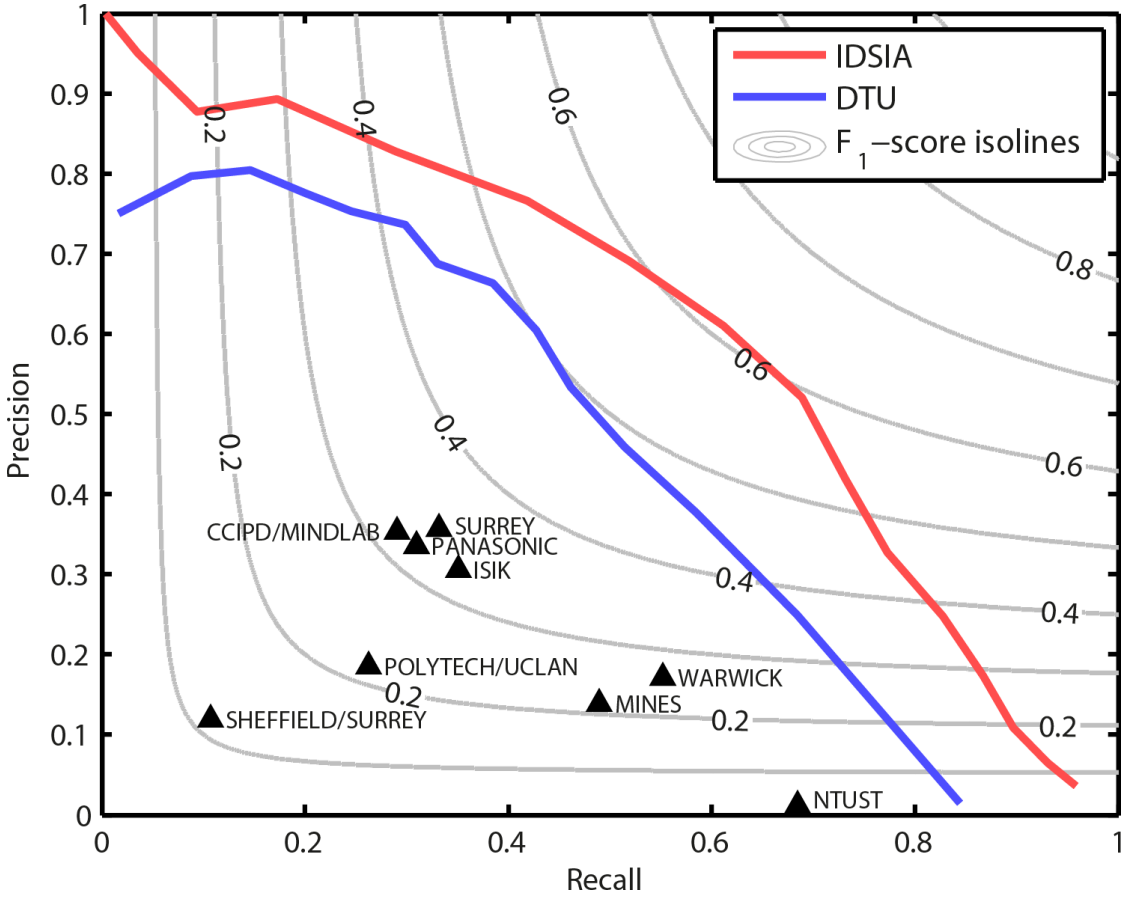
IDSIA
DTU
F1–score isolines
Precision
Recall
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0.2
0.4
0.6
0.8
CCIPD/MINDLAB
SURREY
PANASONIC
ISIK
POLYTECH/UCLAN
WARWICK
MINES
SHEFFIELD/SURREY
NTUST

Reannotation experiment

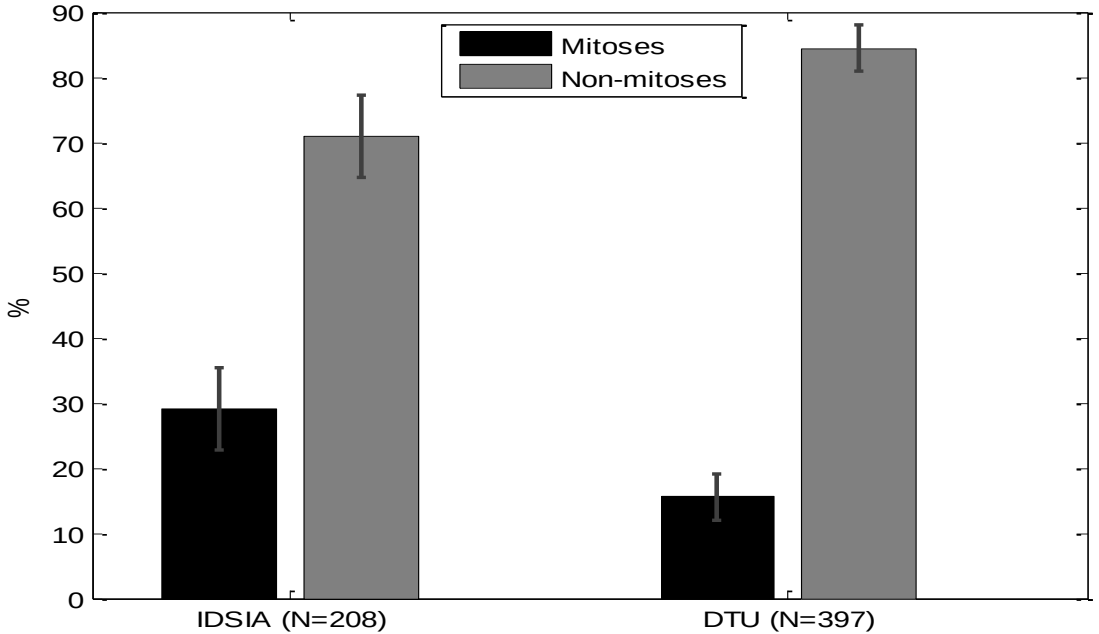

histologists reannotated 30% of all our “False Positives” as actual mitoses they missed during the original annotation

How do you compare with machines?

http://bit.ly/YUYQFG

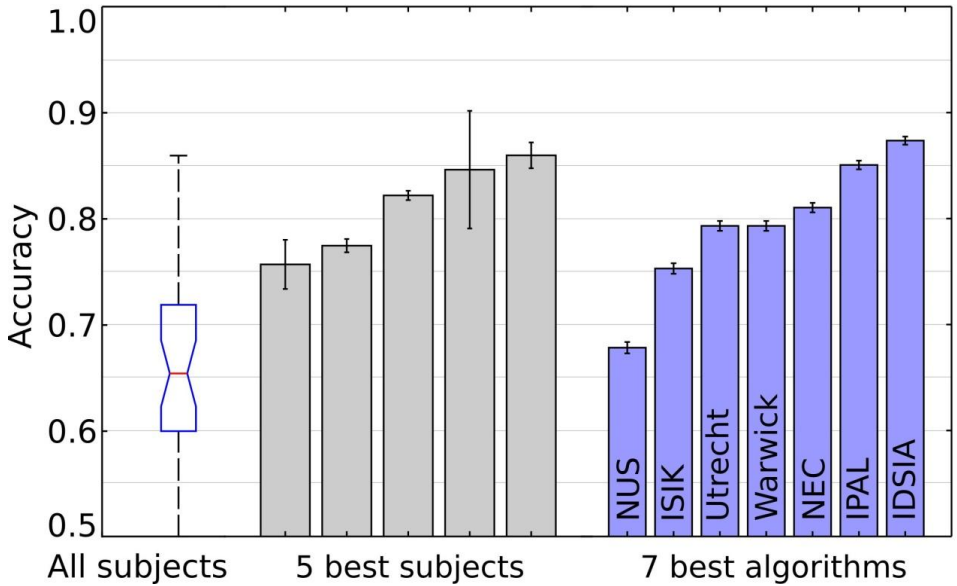

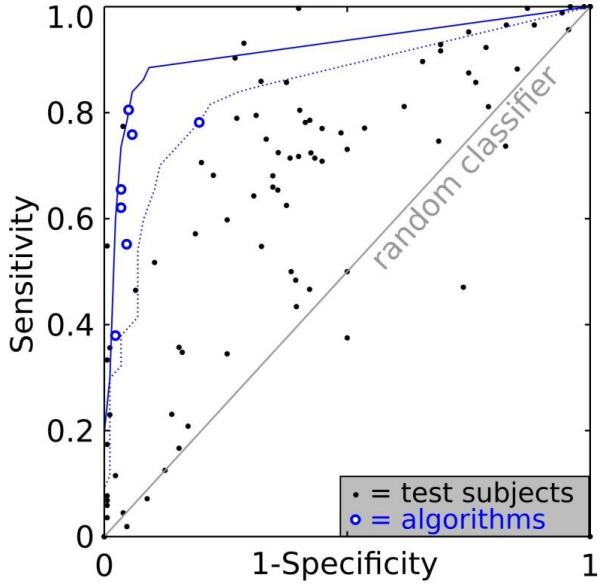

A. Giusti at al. - A Comparison of Algorithms and Humans for Mitosis Detection, ISBI 2014

Results of Mitosis Detection Competitions

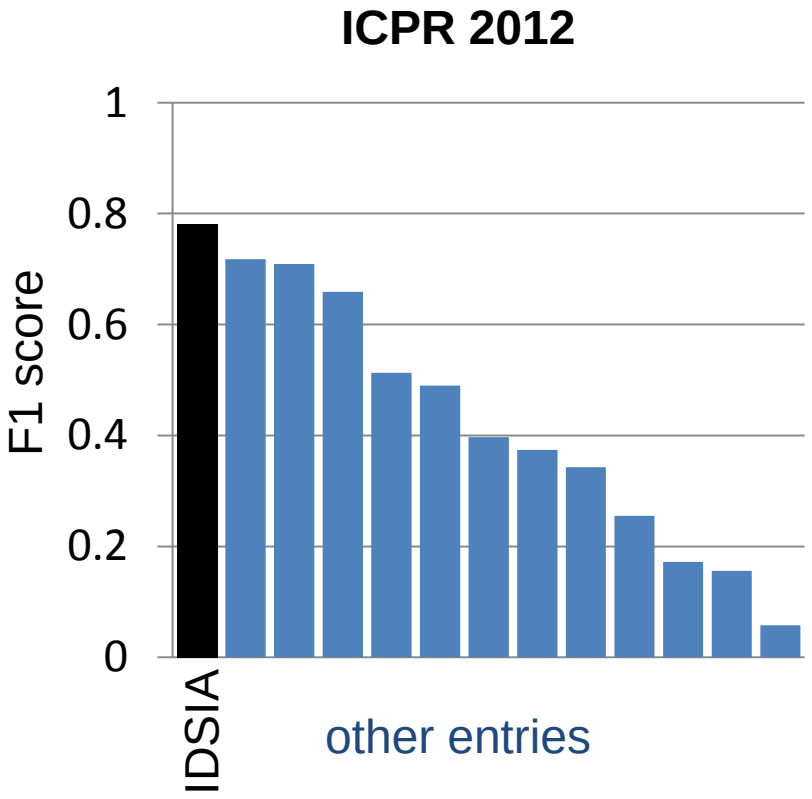

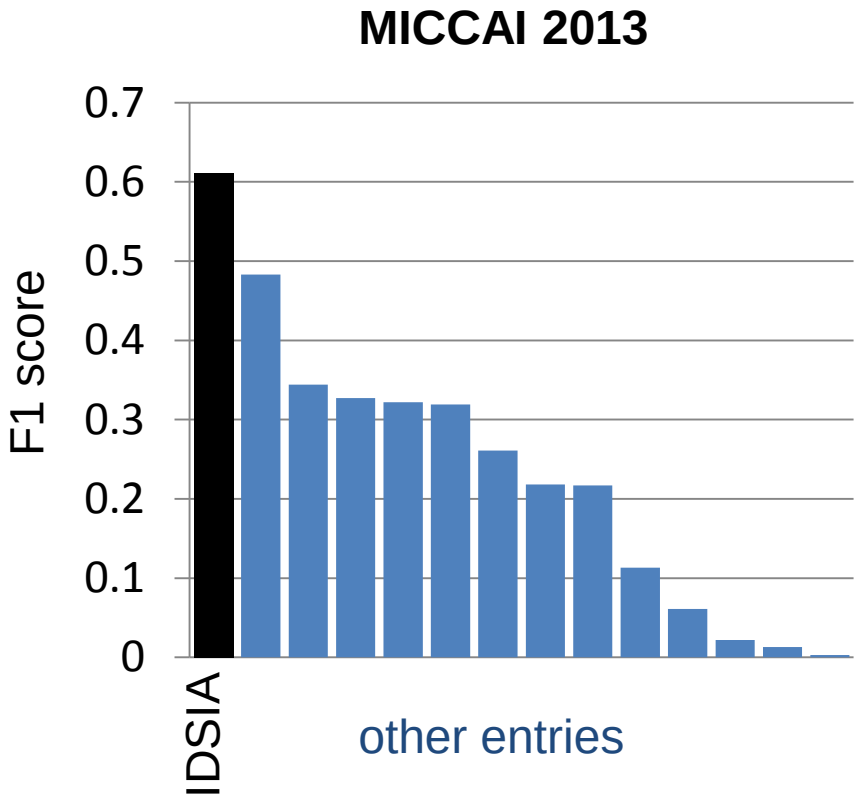

Conclusions

- No need to extract handcrafted features: the network learns powerful features by itself
- Big deep nets combining CNN and other ideas are now state of the art for many image classification, detection and segmentation tasks
- Our DNN won six international competitions
- DNN can be used for various applications: automotive, biomedicine, detection of defects, document processing, image processing, etc.
- **DNN are already better and much faster than humans on many difficult problems**
- GPUs are essential for training DNN. Testing can be done on CPU.
- More info: www.idsia.ch/~ciresan dan.ciresan@gmail.com

Looking for new projects

- Industry
- Academic
  - Unrelated fields: biomedicine, psychology, finance, literature, history
  - Vision for robotics

www.idsia.ch/~ciresan dan.ciresan@gmail.com

Other projects

Neural Networks for Segmenting Neuronal Structures in Electron Microscopy Stacks – ISBI 2012

Training data:

30 labeled 512x512 slices

Test data:

30 unlabeled 512x512 slices

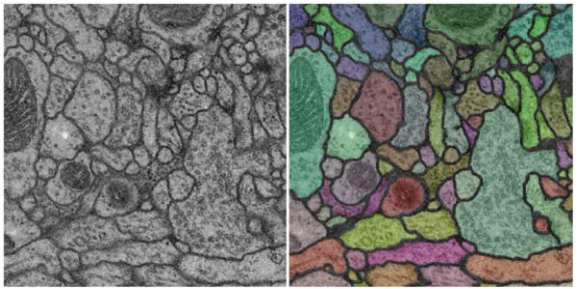

CONNECTOMICS

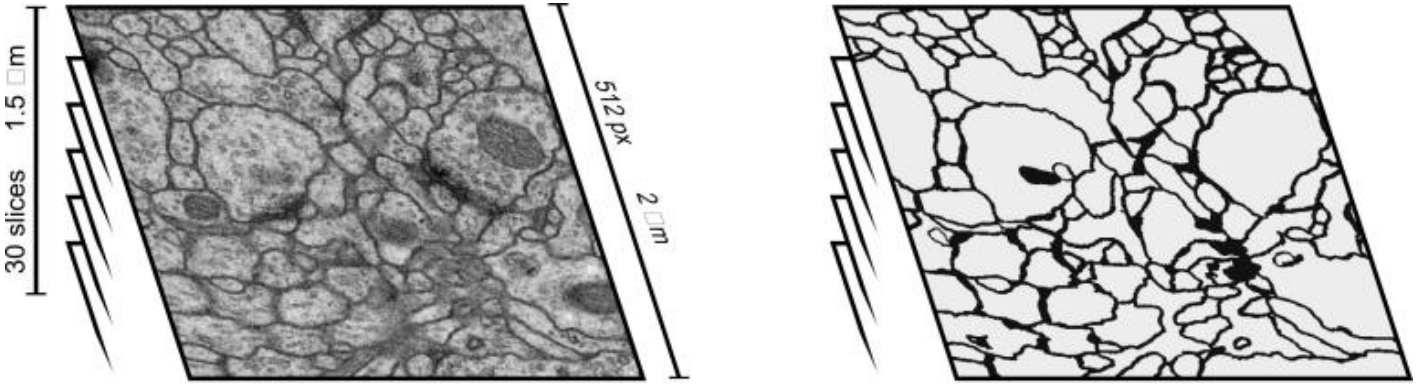

Retina vessel segmentation

- challenging problem
- clinical relevance (e.g. for diagnosing glaucoma)
- state of the art results for DRIVE and STARE datasets
- better than a second human observer

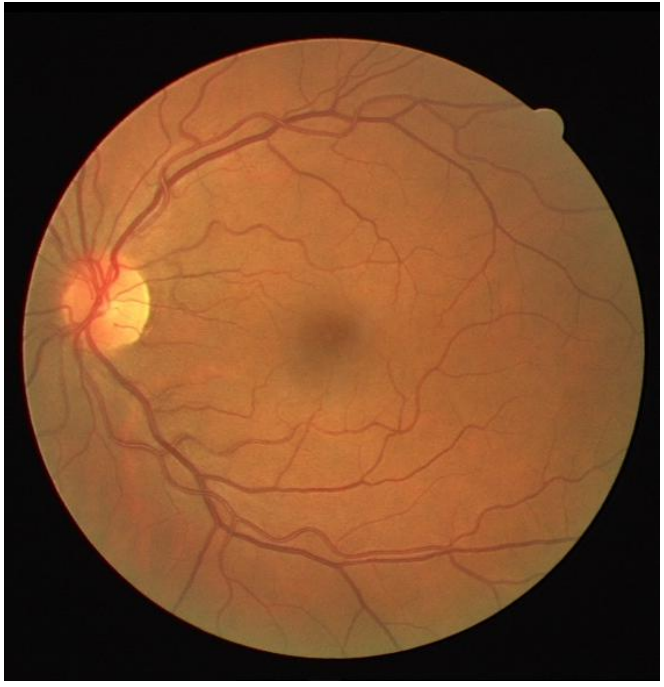

DNN

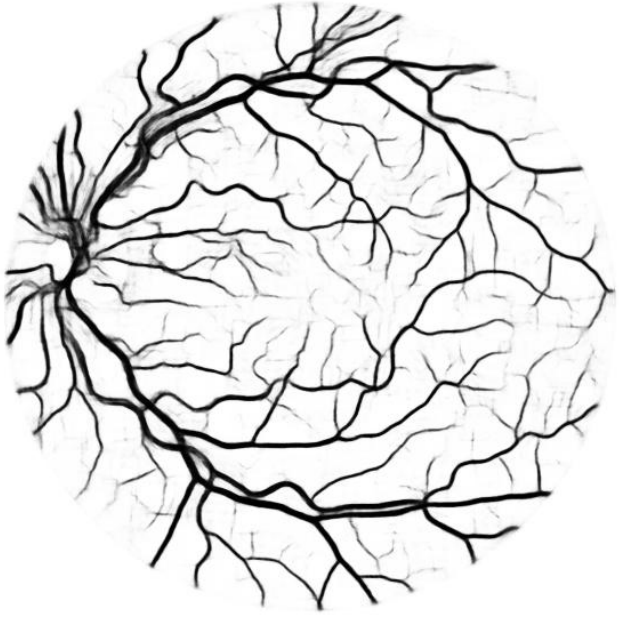

Trail Following Problem

collaboration with Jérôme Guzzi, Alessandro Giusti, Fang-lin He, Juan P. R. Gómez