Perceptor: Under the Microscope with Machine Learning

Camelyon 16 - Challenge Solution

Quincy Wong
qwong77@yahoo.ca

Deep Convolutional Networks are a Serious Game Changer

- AlphaGo
- Reaching or surpassing human levels on Image Recognition Benchmarks: MS COCO, Imagenet
- Enabler for self-driving vehicle technology
- Rapid Advancements
  - Pretrained Networks
  - „End-to-End“ Learning
  - Turn-key Solutions

https://deepmind.com/alpha-go

http://mscoco.org/dataset/#captions-leaderboard

http://image-net.org/challenges/LSVRC/2015/results

http://www.nvidia.com/object/drive-px.html

Background Terminology

- Deep Convolutional (Neural) Networks: a class of supervised machine learning inspired by the biological visual system

- Supervised Learning: Requires a set of training data accompanied with a „Ground Truth“

http://yann.lecun.com/exdb/lenet/

Camelyon16: Eating a „Gigapixel“ Elephant

- Gigapixel: Whole-slide images stored in a multi-resolution pyramid structure
- 500MB - 4GB single Image Slide
- Contains “uninteresting“ empty space

<u>Approach:</u>

- Focus on a single magnification level (20x)
- Tumor slides are the most interesting

Curating the Training Set

- Ideally the training set / images should contain:
  - Image regions containing normal and tumor cells
  - Different scanners and staining variations
  - Different appearance of tumors
  - Shapes in which the tumors manifest around normal cells

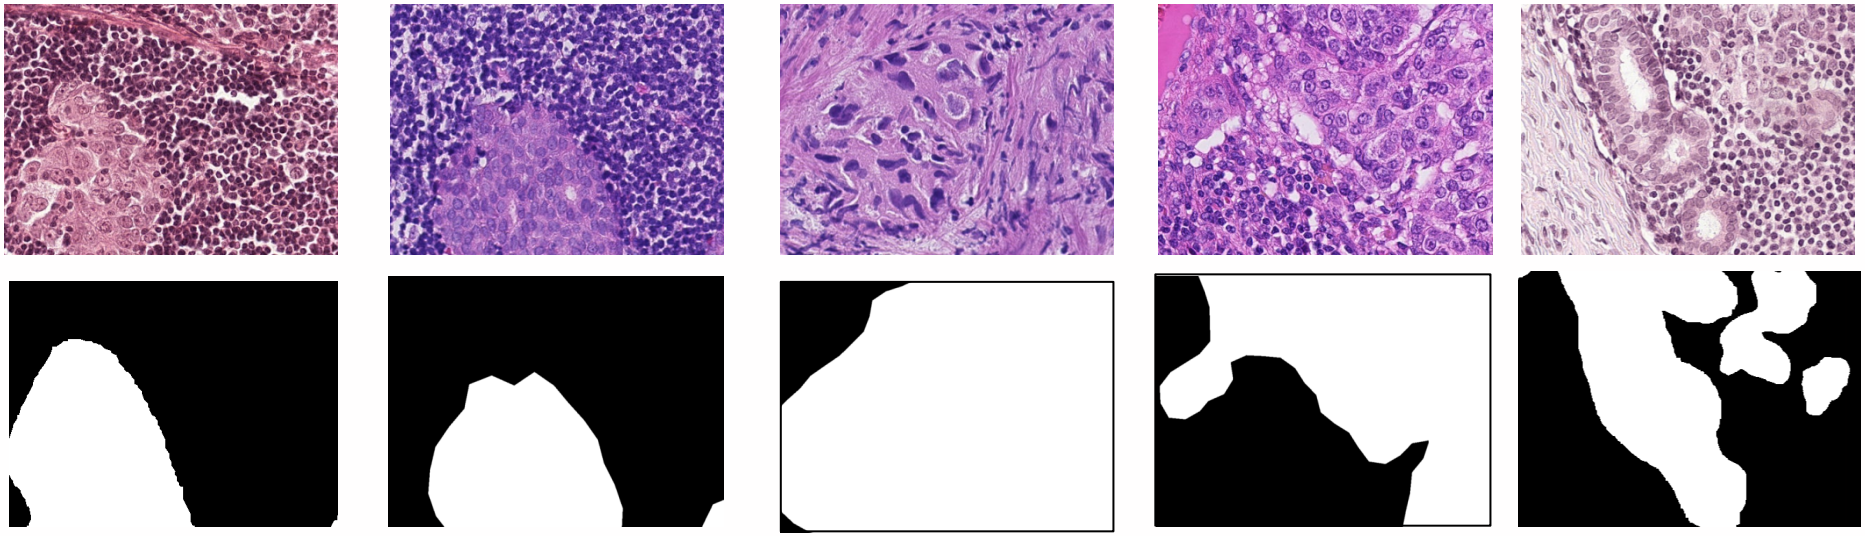

Choosing a Technology Stack for Rapid Development

- Python Libraries
  - Scikit-image, SciPy, Pandas, Matplotlib, Openslide
- Caffe SegNet
- Docker Containers with GPU passthrough
- .. and more Python to bring them all, and in the darkness bind them

Deep Convolutional Networks: Heart of the Solution

- Caffe SegNet: deep learning framework
  - http://caffe.berkeleyvision.org/
  - http://arxiv.org/abs/1511.00561
- Allows for Pixel Level Classification
- 37 layer, fully convolutional architecture
- Pretrained VGG-16 model
- Trained 40-60 Thousand Epochs
- Approx. 2 days using Nvidia GTX 960

Results

Training Eg. 1 (ke = thousand epochs)

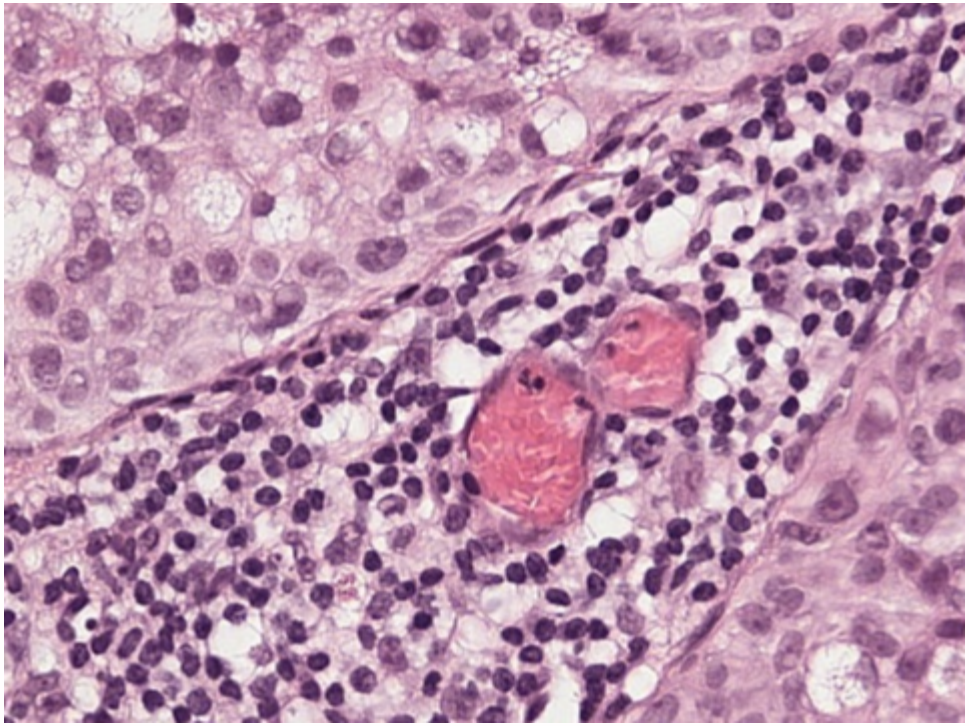

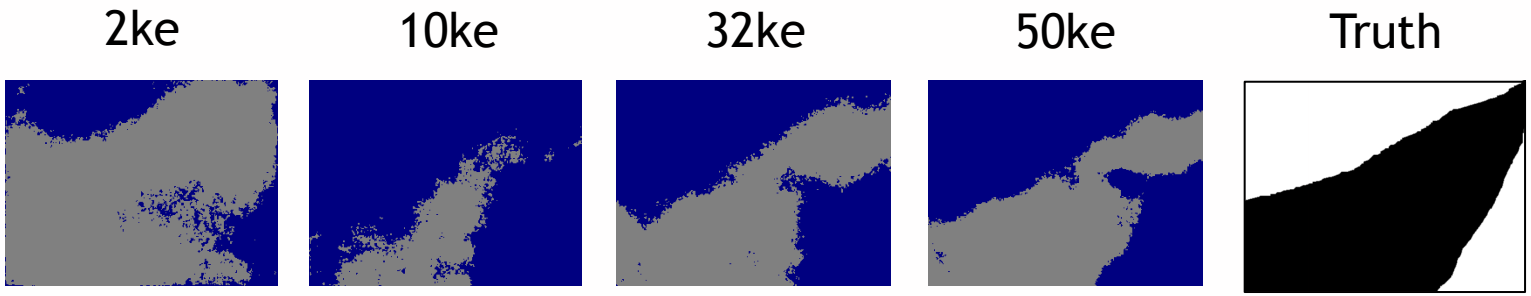

Training Eg. 2 (ke = thousand epochs)

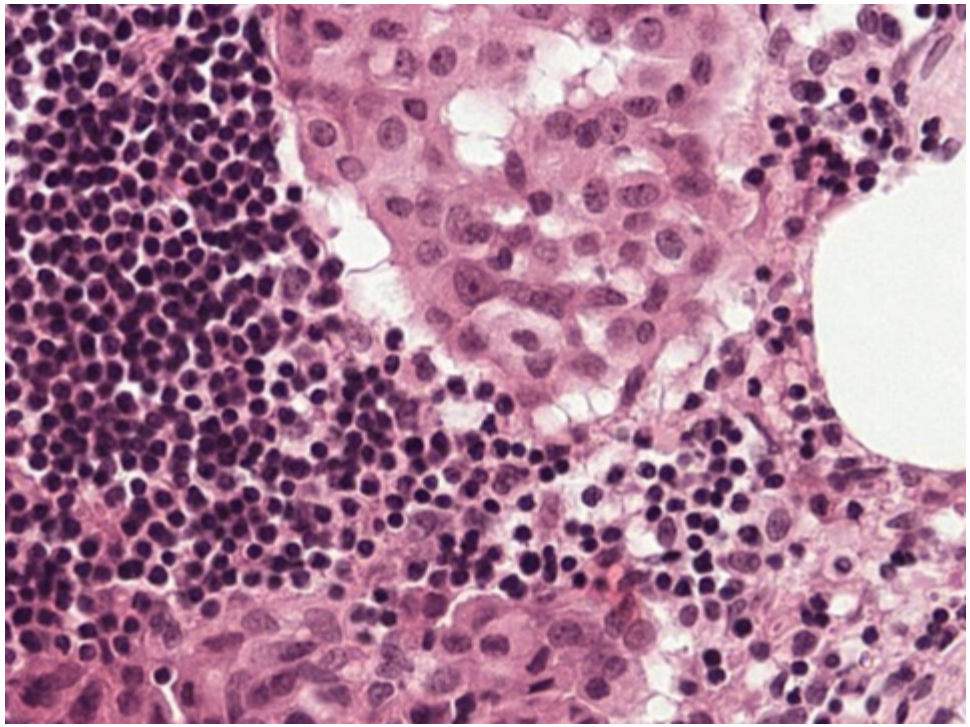

2ke

10ke

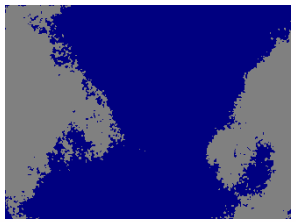

32ke

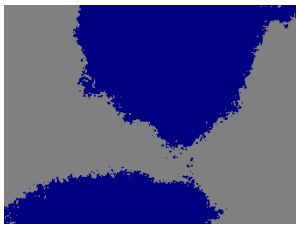

50ke

Truth

Major Axis Measurements Eg. 1

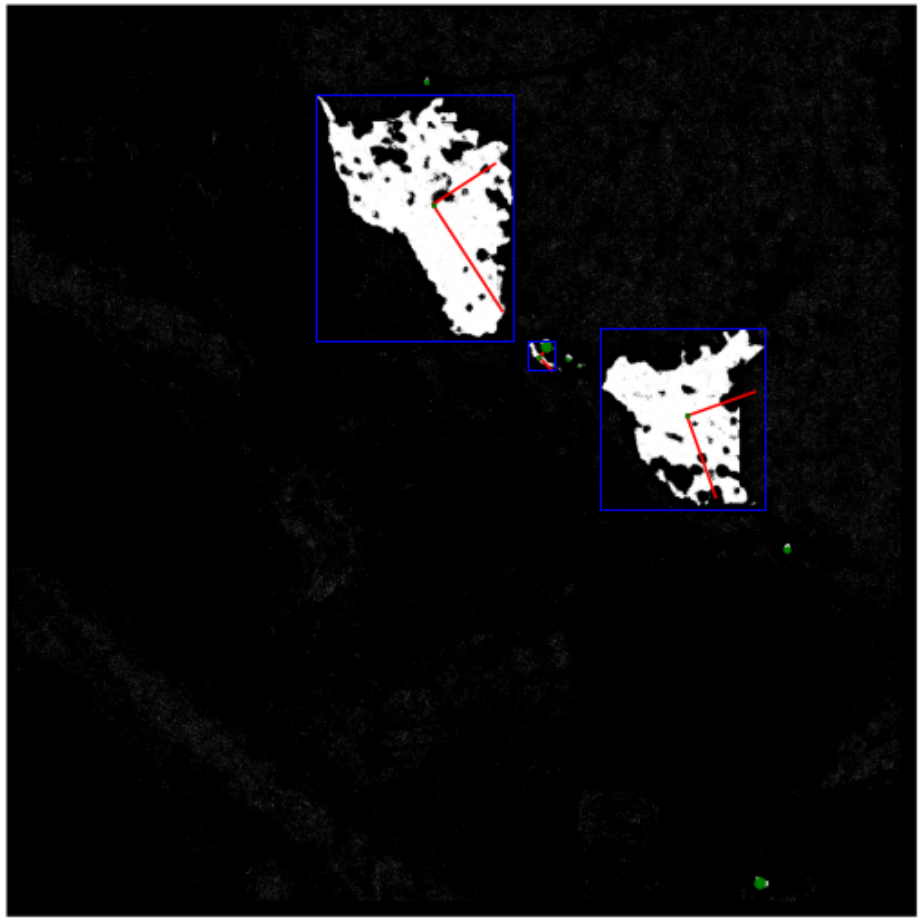

- Isolated tumor cells (ITC)
- Micro / Macro metastases

Major Axis Measurements Eg. 2

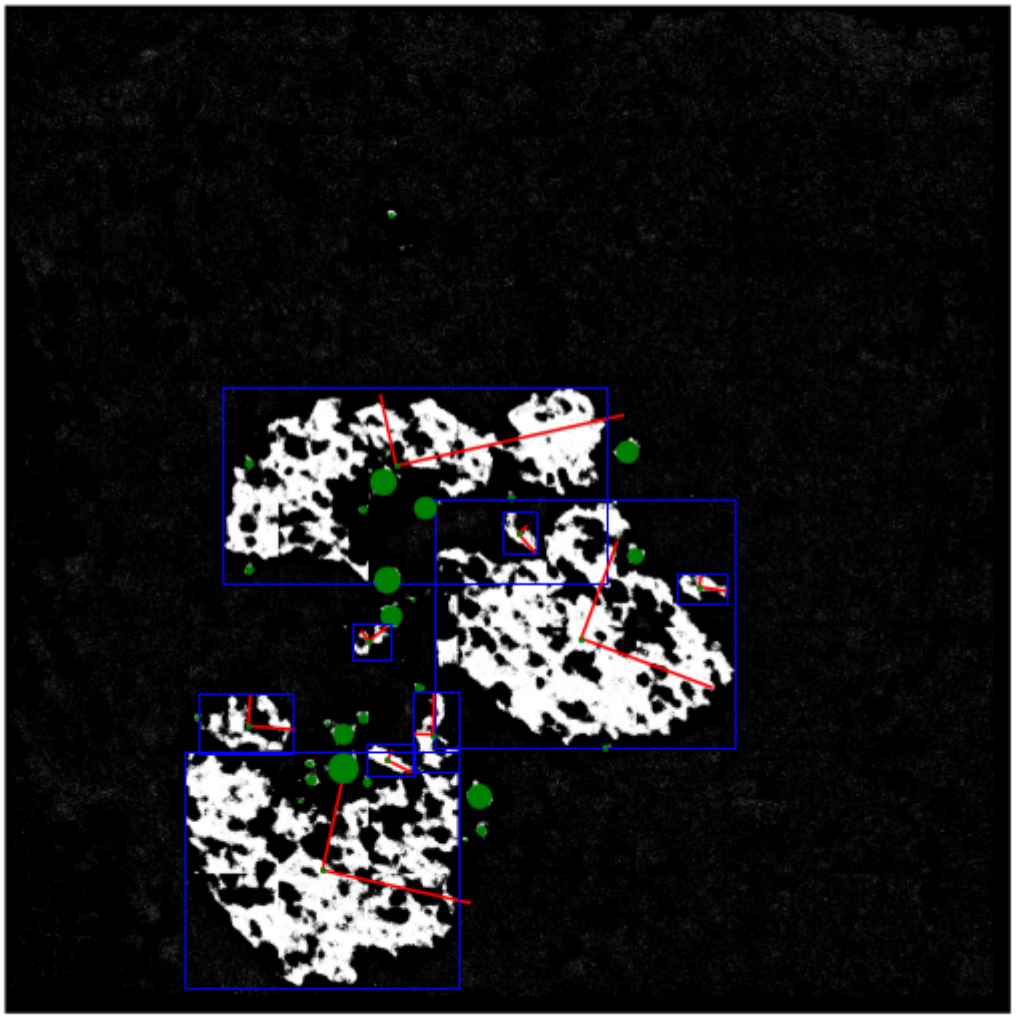

- Isolated tumor cells (ITC)
- Micro / Macro metastases

Major Axis Measurements Eg. 3

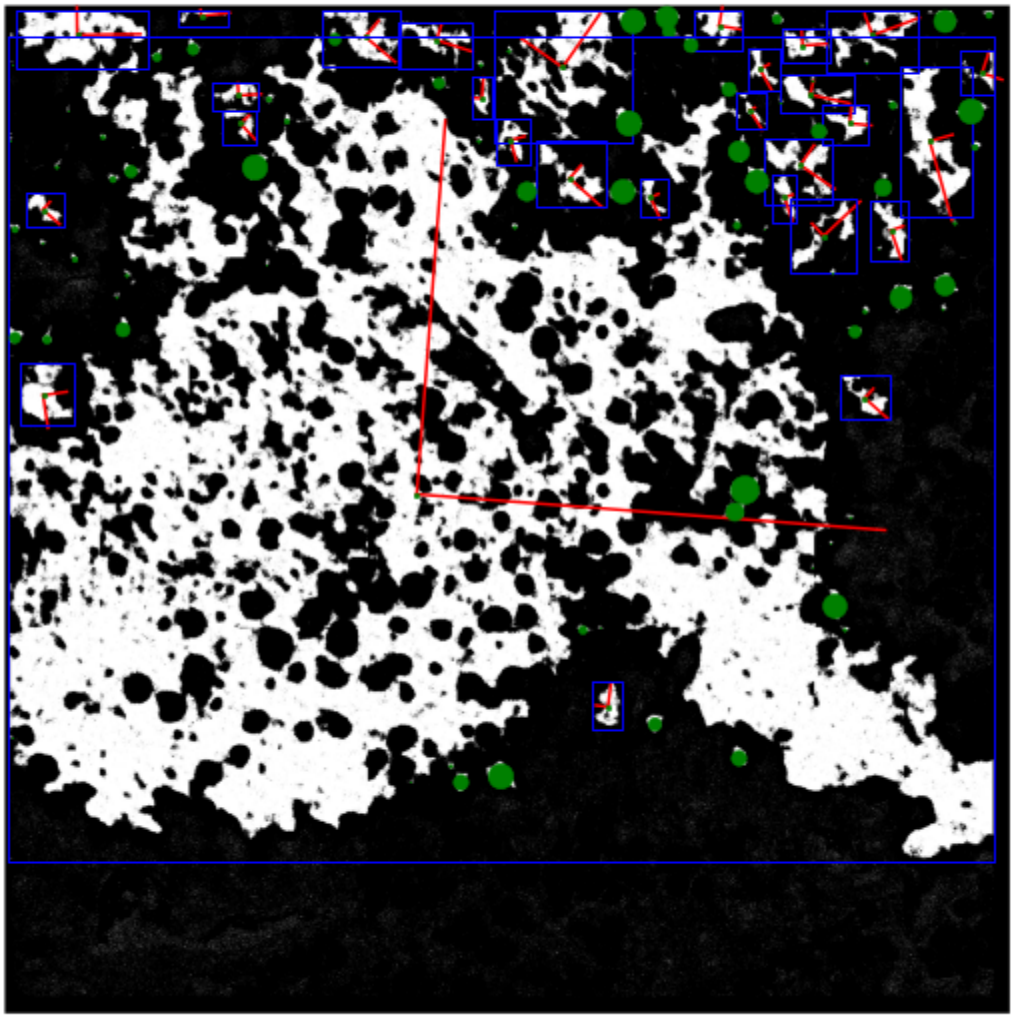

- Isolated tumor cells (ITC)
- Micro / Macro metastases

Take-Aways

- The Impact of Machine Learning is real
  (Threat to humanity, probably not yet)
- Domain Knowledge, Feature Engineering marginalized through End-to-End learning
- Enables small teams (or individuals) to seriously punch above their weight when leveraged