Self-driving cars

Using **deep learning**
Optimization for the best performance

Marek Bardoński

Short introduction

- Working as a software research engineer at NVIDIA on TensorRT - a tool used for improving the performance of inference.

- Helping Google, Facebook, Tesla, Daimler and other to make their networks as fast as possible.

- Currently on vacation, traveling between multiple countries in Asia and giving lectures.

Where deep learning in cars?

- Automatic cruise control
- Lane keeping assistant
- Driver status monitoring, eye tracking
- Anti-collision systems
- *Self-driving*

When we will have self-driving cars, that not require driver?

- Tesla ~ 1 year

- Daimler, Toyota, NVIDIA, Google, others ~ 5 years

What kind of algorithms are used for self-driving?

- **Region Proposals** for detection
  Trends: Customized faster RCNN

- **Convolutional neural networks** for classification
  Trends: Customized Inception V4, increasing depth, custom layers (eg. leaky pooling)

- **Reinforcement learning** for driving policy

How CNNs are optimized?

- DrivePX2 system that have two Tegra GPUs (DriveTX2 for drones)

- TensorRT tool optimize the network - layers merging, kernel selection, reducing precision (FP16, INT8).

- CuDNN (or TensorRT) execute GPU-assembly kernels to process each layer.

How does my job looks like?

- Talking with customers
- Measuring the performance of networks
- Assisting customers with model selection
- Implementing CUDA kernels
- Implementing GPU-assembly kernels

Medical diagnostic

Using **deep learning**
on a cancer recognition problem example

How deep learning is affecting healthcare?

- IBM Watson for Oncology is used for recommending personalized, evidence-based cancer care options - adopted already in Manipal Hospitals in India as a second opinion provider.

- Convolutional Neural Networks are used for recognizing cancer on mammography images - being deployed by Enlitic in Australia.

Cancer

- Approximately 39.6% of men and women will be diagnosed with cancer.
- One of the main reasons of death.
- Usually diagnosed by radiologist by looking at MRI / X-ray images.

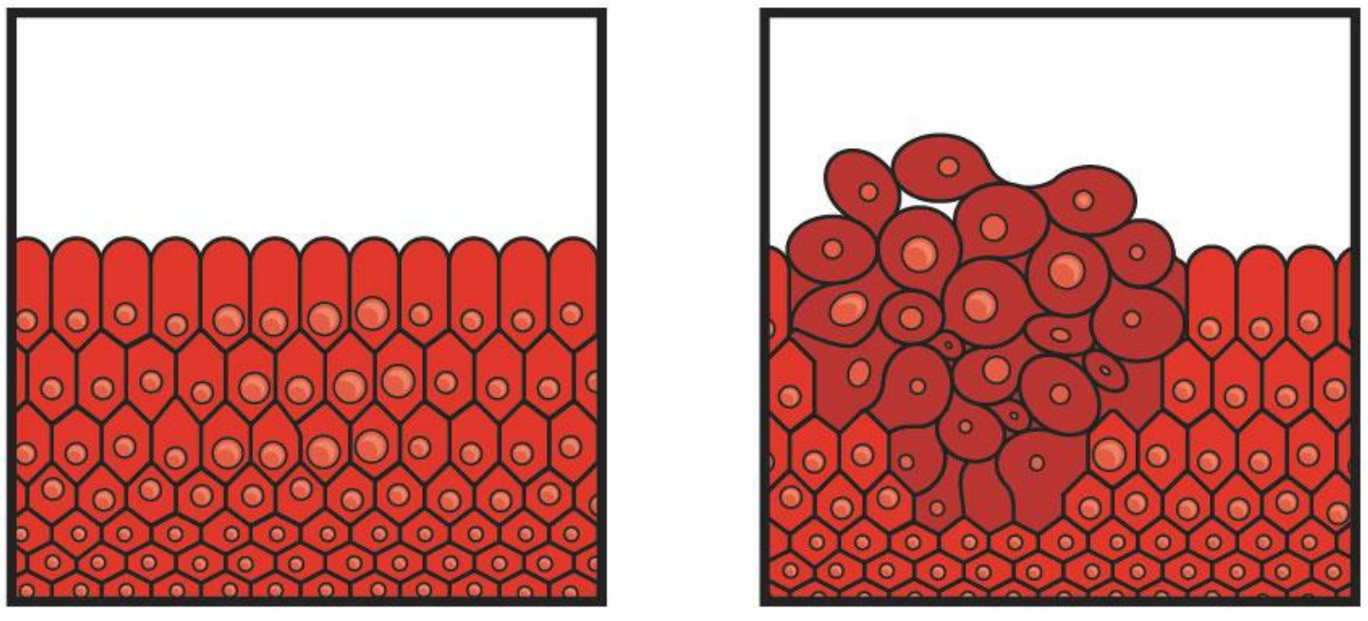

Normal cells

Cells forming a tumour

Breast cancer

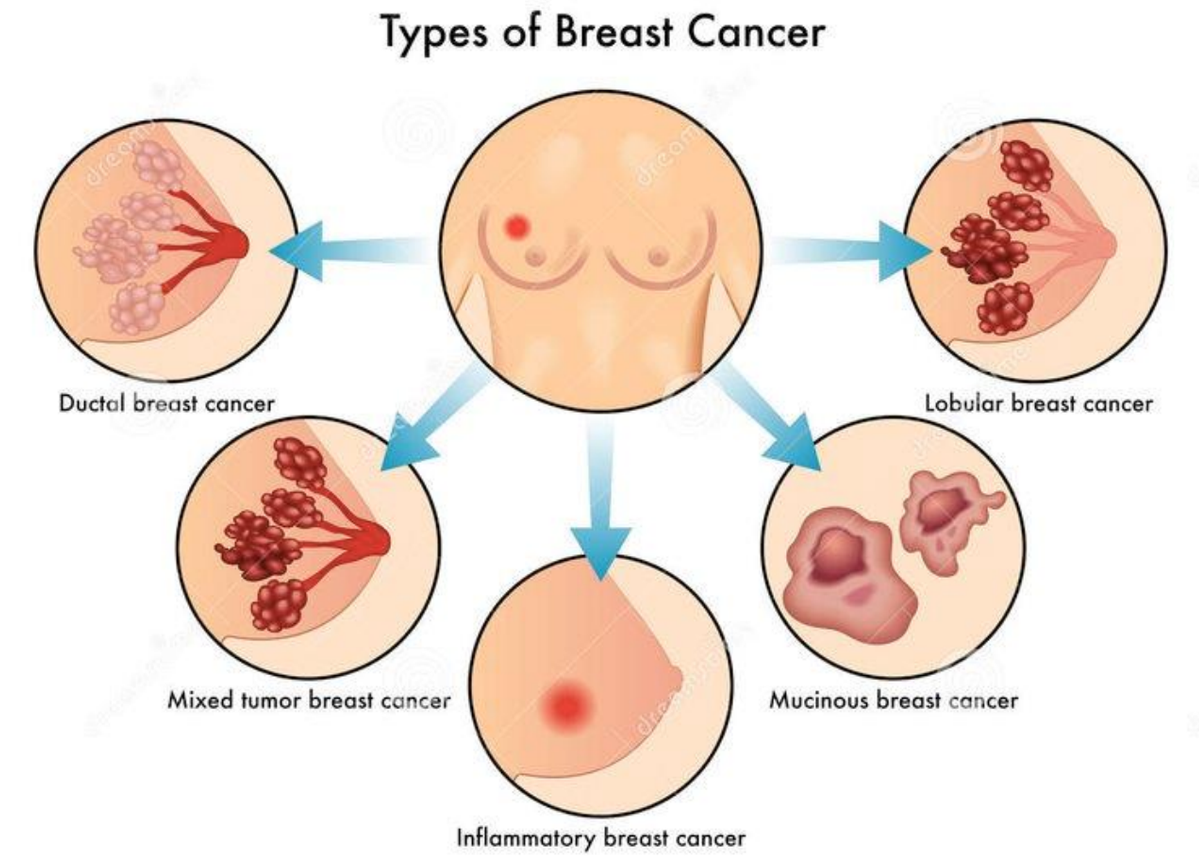

A chance that woman will die from a breast cancer is **2.7%**

Main problems

- Mammography screening reduced the overall mortality **by 20% only**

- When diagnosing cancer, US radiologist recall is around **0.84**

- Women that start annual screening at 40 years old, have a **7%** cumulative risk of a false-positive biopsy over 10 years.

- Scientists are debating if the mammography screening should be discontinued.

How good is deep learning?

The algorithm developed by Yaroslav Nikulin have **0.8** specificity (true negative rate) at **0.8** sensitivity (true positive rate), when predicting if the person will have been confirmed with cancer during next 12 months, basing on a set of two standard mammography images.

The current performance of US radiologist is **0.91** specificity at **0.84** sensitivity (2004-2008) — but **it's much lower in 3-rd world countries.**

How does his model looks like?

General outline: inspired by VGG, FCN, ResNet

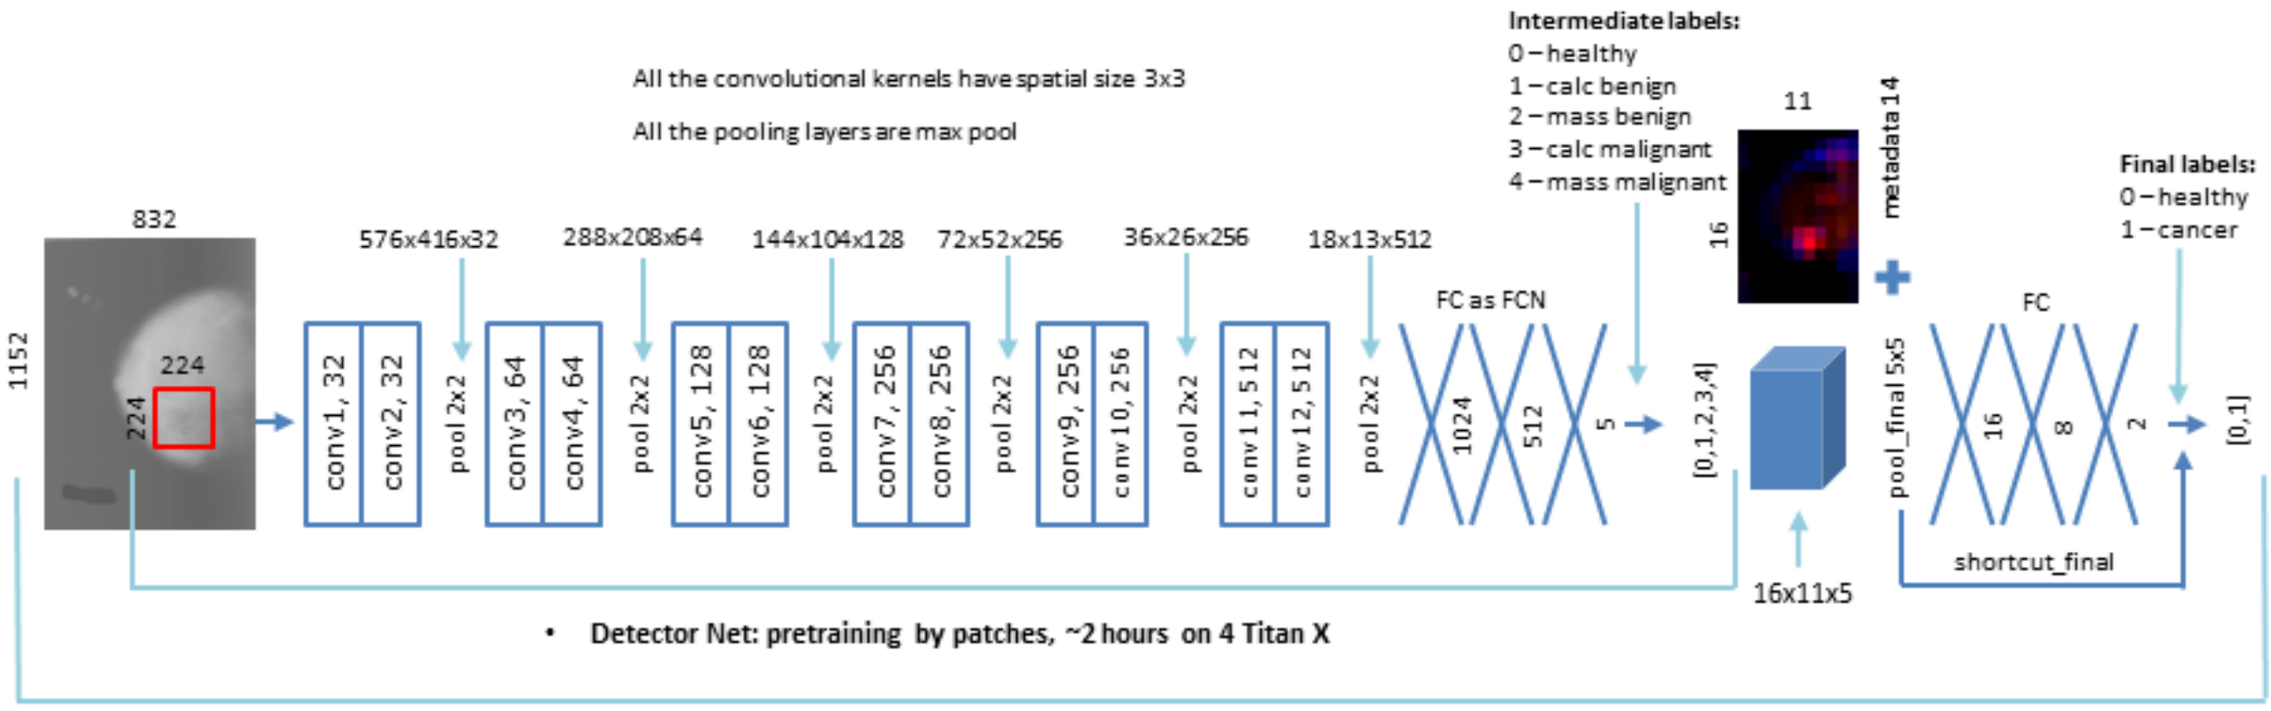

- Detector Net: pretraining by patches, ~2 hours on 4 Titan X
- End-to-end finetuning by images, ~20 hours on 4 Titan X

Data augmentation:
Flips: hor & ver
Rotation: ±20°
Zoom: ±20%
Shear: ±20%
Channel Shift: ±10%

In the cloud:
Finetune with Adam and EMA = 0.99
Restore EMA of the best snapshot
Finetune with SGD lr=10^-4
Pick the best one!

During the inference:
Images per breast: x2 (x3, x4...)
Flips: hor & ver -> x4
Network ensemble: x4
Average everything
PROFIT!

Customization techniques used

- Squeezing of the model. Reducing the width while keeping the height
- Additional pooling layer at the end of detector net.
- Smoothing the bottleneck created by few output neurons by reducing neurons in the second fully connected layer
- Adding batch normalization

Questions!

References

- https://www.ncbi.nlm.nih.gov/pubmed/26501537
- https://www.synapse.org/#!Synapse:syn4224222/wiki/401748
- http://www.bcsc-research.org/statistics/performance/screening/2009/perf_age.html