



LUXOFT UKRAINE

Computer Vision – Practical Case Studies

Ihor Starepravo

«AI Ukraine»

Conference 9 of October, 2016

LXFT
LISTED
NYSE





Please don't shoot the piano player!



He's doing the best
that he can!



LUXOFT OVERVIEW. COMPANY FACTS



11,000+

Employees Worldwide



150+

Active Clients,
30% from Fortune 500

Industry Verticals

Financial Services

Trading & Risk Platforms
Regulatory
Reference Data Platforms
Digital

Automotive

Autonomous driving
HMI
Connected Cars
Under the hood

Telecom

Intelligent NFV

Travel and Aviation


eBooking
eCRM

Technology

Connected devices
Smart mobility
Information security

Energy

Energy grids



3

ITO geographies
(CEE, LATAM, APAC)

5

Continents

17

Countries

33

Offices

Why LUXOFT is Giving This Talk?

Global Software Engineering:

- Low-Power GPU Software
- Custom Vision Software
- Low System Optimization

A.I. & Computer Vision

- Looking for talents and partners
- 5+ years in CNN/DNN
- 10+ vision projects:
 - ADAS, Security, Enterprise, Retail, Pharma



PROJECTS IN THIS TALK

CAFFE on ARM Mali




ADAS: HOG/BLINK +
DNN Pedestrians Tracking




Spatial-temporal analysis:
Pharma Surveillance



OTHERS




DEMAND PYRAMID



CAFFE at ARM Mali T880MP12

- ♦ **Problem:** achieve highest values in forward operation for any DNN application
- ♦ **Solution:** AlexNet@CAFFE image visual analysis in real-time
- ♦ **DNN:** ImageNet 1000 tags, 8 layers, ImageNet, Top-5 error, %: 16,4



ARM Mali Accelerated CAFFE

Original. All FPS measured on Galaxy S7:

- Run existing DNN framework: CAFFE
- =0.7 FPS (EIGEN OpenCL library)

Open Source CPU, 1 thread	Open Source GPU OpenCL (ViennaCL)	Open Source CPU multithreaded, NEON	LUXOFT
0.7 FPS	1.2 FPS	2.5 FPS	5.4 FPS

CPU Optimization (not a through road):


- Optimized version for Android: DNN optimized OpenBLAS: =1.2 FPS

OpenMP and NEON → +2

GPU Optimizations:

- Better OpenCL implementation on ViennaCL library → +0.5 FPS
- Found bottleneck: SGEMM functions → >85% CPU/GPU
 - → ARM F16-SGEMM (workgroup size, vectorization, etc) → x2 FPS (No F16 support in CAFFE)
 - → Rewrite SGEMM (workgroup size, vectorization, etc) → +4.5 FPS


Final optimized performance: 5-6 FPS



DREAM2020: AUGMENTED VISION NAVIGATION



DNN for ADAS



◆ Step 1: HOG (Histogram Oriented Gradient)


- OpenCV implementation → 100 FPS false-positive ratio: → 40%
- NOT suitable for Real-World

◆ Step 2: HOG + DNN

- AlexNet: 6 layers, 30 classes
- AMD Tahiti XT/Hawaii → 1.3 FPS
- false-positive: → 0%, false-negative: → 10%

◆ Step 3: Faster-RCNN (work in progress)

- AMD Tahiti XT/Hawaii: → 0,3 FPS, RAM: 2GB
- NVIDIA GTX 1080: → 9 FPS, RAM: 2GB
- Solution: SGEM -> FFT




Spatial-temporal analysis: Pharma Surveillance

- ◆ **Problem:** 24/4x365 laboratory tests generates massive amount of video data, that is difficult to maintain, hard to analyze by humans and provides no scientific value for the client. Video data is not integrated within the Laboratory Data Space.
- ◆ **Solution:** DNN backed specialized video analysis tool that provides reliability and accuracy of video cognition adding new type of the 'sensor' to data scientists.
- ◆ **DNN:** Pipeline of YOLO + DNN + DT



YOLO+DNN+DT

- ◆ Vision Trainer:
 - Train-set preparation: classifications, rects
 - Quick forward-test
- ◆ YOLO
 - Instrument elements mapping
 - 45 hours training time → [NVIDIA GTX TitanX, 12GB](#)
- ◆ DNN
 - AlexNet based: 6 layers,
 - Training set: 4 classes, 1000+ per class
 - 45 min training time → [NVIDIA GTX TitanX, 12GB](#)
- ◆ Dense trajectory features: motion analysis
 - Dense trajectory features [Wang, 13]



Conclusions: Pros & Cons



- Highly flexible
- Adaptable to the new tasks
- Great for complex noisy data
- Deterministic latency



- Undebuggable
- Compute & power & memory intensive
- Lacks support for video
- Not quite understood by developers

Conclusions



- ◆ Deep Learning is in it's infancy
- ◆ Development process insights:
 - Be ready for chaotic tools landscape: HW, Tools, Nets, etc.
 - Quality of training set = application quality
 - Training set prep → up to 80% time
- ◆ Don't break up with traditional algorithmic approach!



THANK YOU!



Ihor Starepravo



Director, Embedded Systems
IStarepravo@luxoft.com