INTEL® DISTRIBUTION OF OPENVINO™ TOOLKIT

Accelerate Computer Vision with OpenVINO™ toolkit

(Open Visual Inference & Neural Network Optimization) - Formerly Intel® Computer Vision SDK

What it is

A toolkit to fast-track development of **high performance computer vision** and **deep learning into vision applications**. It enables deep learning on hardware accelerators and easy **heterogeneous** execution across Intel® platforms. Components include:

- Intel® Deep Learning Deployment Toolkit (model optimizer, inference engine)
- Optimized functions for OpenCV*, media encode/decode, and more

Why important

Demand is growing for intelligent vision solutions. **Deep learning revenue** is estimated to grow from $655M in 2016 to **$35B by 2025**[1]. This requires **developer tools** to integrate computer vision, deep learning, and analytics processing capabilities into applications, so they can help **turn data into insights that fuel artificial intelligence**.

Users

Software developers, **data scientists** working on vision solutions for surveillance, robotics, healthcare, office automation, transportation, & more.

OpenVINO™ version is 2018 R2

[1]Tractica 2Q 2017

Free Download

software.intel.com/openvino-toolkit

What's Inside the OpenVINO™ toolkit

Intel® Deep Learning Deployment Toolkit

Model Optimizer
Convert & Optimize

IR

Inference Engine
Optimized Inference

20 Pre-trained Models & Code Samples

IR = Intermediate Representation file

Traditional Computer Vision Tools & Libraries

Optimized Libraries

OpenCV*

OpenVX*

Photography Vision

Code Samples

For Intel® CPU & CPU with integrated graphics

Increase Media/Video/Graphics Performance

Intel® Media SDK
Open Source version

OpenCL™ Drivers & Runtimes

For CPU with integrated graphics

Optimize Intel® FPGA

FPGA RunTime Environment
(from Intel® FPGA SDK for OpenCL™)

Bitstreams

FPGA – Linux* only

OS Support CentOS* 7.4 (64 bit) Ubuntu* 16.04.3 LTS (64 bit) Microsoft Windows* 10 (64 bit) Yocto Project* version Poky Jethro v2.0.3 (64 bit)

Intel® Architecture-Based Platforms Support

Benefits of the OpenVINO™ toolkit

Harness the Power of Intel® Processors: CPU, CPU with Integrated Graphics, FPGA,VPU

ACCELERATE PERFORMANCE

Access Intel computer vision accelerators.
Speed code performance.
Supports heterogeneous processing
& asynchronous execution.

INTEGRATE DEEP LEARNING

Unleash convolutional neural network (CNN) based deep learning inference
using a common API & 20 pre-trained models.

Up to 19.9x increase[1]

SPEED DEVELOPMENT

Reduce time using a library of optimized OpenCV* & OpenVX* functions, & 15+ samples.

Develop once, deploy for current
& future Intel-based devices.

INNOVATE & CUSTOMIZE

Use the increasing repository
of OpenCL™ starting points in OpenCV*
to add your own unique code.

[1]Performance increase comparing certain standard framework models vs. Intel-optimized models in the Intel® Deep Learning Deployment Toolkit. See Benchmarks slides for details.
Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

**Certain technical specifications and select processors/skus apply. See product site for more details.
OpenVX and the OpenVX logo are trademarks of the Khronos Group Inc.
OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos

Intel as Best Choice for Vision Solutions

Deliver High Performance Computer Vision & Deep Learning – Transform Data & Results into Artificial Intelligence

Intel offers the broadest portfolio of hardware and software that help you

- Accelerate workloads for a wide range of solutions and vertical use cases
- Increase application performance through Intel accelerators and flexible heterogeneous architectures[1] (CPU, CPU w/integrated graphics, FPGA, and Vision Processing Units (VPU))
- Drive power, cost and development efficiencies to designs and applications for cameras, gateways, network video recorders (NVR), and servers
- Enable deep learning capabilities for smarter, faster analytics – transform data into artificial intelligence for competitive advantage

[1]Certain technical specifications and select processors/skus apply. See product site for more details.

Increase Deep Learning Workload Performance on Public Models using OpenVINO™ toolkit & Intel® Architecture

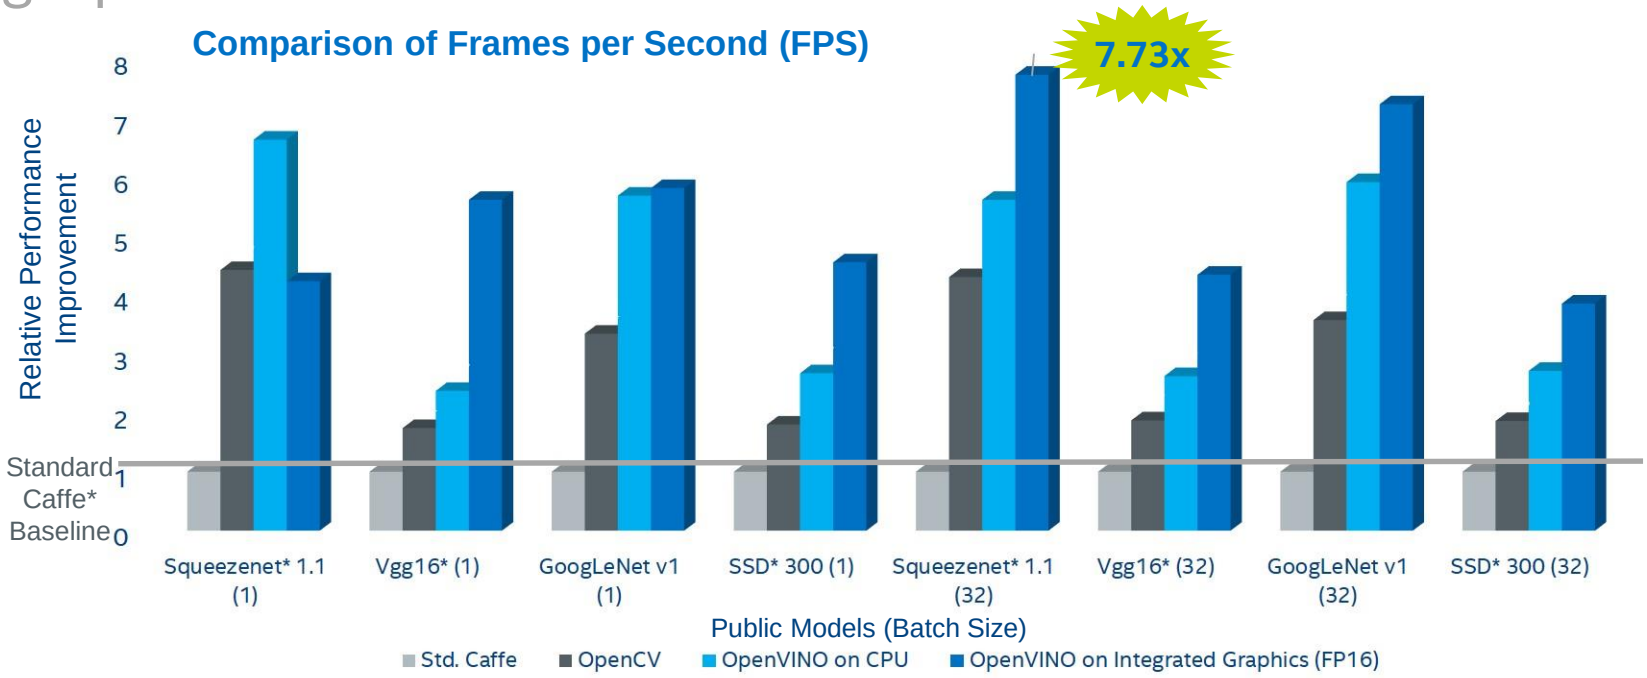

Fast Results on Intel Hardware, even before using Accelerators

[1]Depending on workload, quality/resolution for FP16 may be marginally impacted. A performance/quality tradeoff from FP32 to FP16 can affect accuracy; customers are encouraged to experiment to find what works best for their situation. The benchmark results reported in this deck may need to be revised as additional testing is conducted. The results depend on the specific platform configurations and workloads utilized in the testing, and may not be applicable to any particular user's components, computer system or workloads. The results are not necessarily representative of other benchmarks and other benchmark results may show greater or lesser impact from mitigations. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks. **Configuration:** Intel® Core™ i7-6700K CPU @ 2.90GHz fixed, GPU GT2 @ 1.00GHz fixed Internal ONLY testing, performed 4/10/2018 Test v312.30 – Ubuntu* 16.04, OpenVINO™ 2018 RC4. Tests were based on various parameters such as model used (these are public), batch size, and other factors. Different models can be accelerated with different Intel hardware solutions, yet use the same Intel software tools. Benchmark Source: Intel Corporation.

Increase Deep Learning Workload Performance on Public Models using OpenVINO™ toolkit & Intel® Architecture

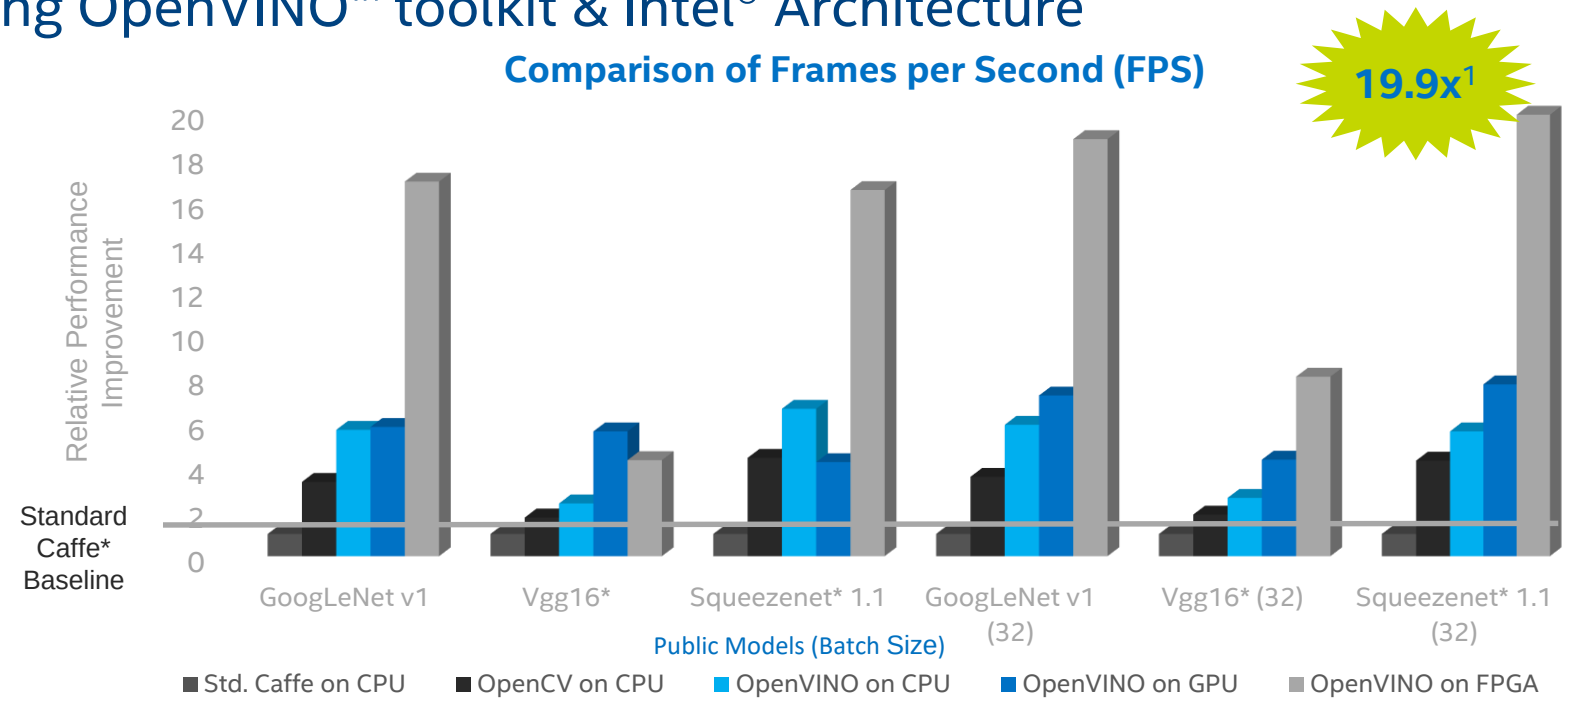

Get an even Bigger Performance Boost with Intel® FPGA

[1]Depending on workload, quality/resolution for FP16 may be marginally impacted. A performance/quality tradeoff from FP32 to FP16 can affect accuracy; customers are encouraged to experiment to find what works best for their situation. The benchmark results reported in this deck may need to be revised as additional testing is conducted. The results depend on the specific platform configurations and workloads utilized in the testing, and may not be applicable to any particular user's components, computer system or workloads. The results are not necessarily representative of other benchmarks and other benchmark results may show greater or lesser impact from mitigations. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks. **Configuration:** Intel® Core™ i7-6700K CPU @ 2.90GHz fixed, GPU GT2 @ 1.00GHz fixed Internal ONLY testing, performed 4/10/2018 Test v312.30 – Ubuntu* 16.04, OpenVINO™ 2018 RC4. Intel® Arria 10-1150GX FPGA. Tests were based on various parameters such as model used (these are public), batch size, and other factors. Different models can be accelerated with different Intel hardware solutions, yet use the same Intel software tools. Benchmark Source: Intel Corporation.

Call to Action, Resources

Download ▸
Free OPENVINO™ toolkit

Learn More

Get started quickly with:

- Developer resources
- Intel® Tech.Decoded online webinars, tool how-tos & quick tips
- Hands-on in-person events

Support

- Connect with Intel engineers & computer vision experts at the public Community Forum

Select Intel customers may contact their Intel representative for issues beyond forum support.

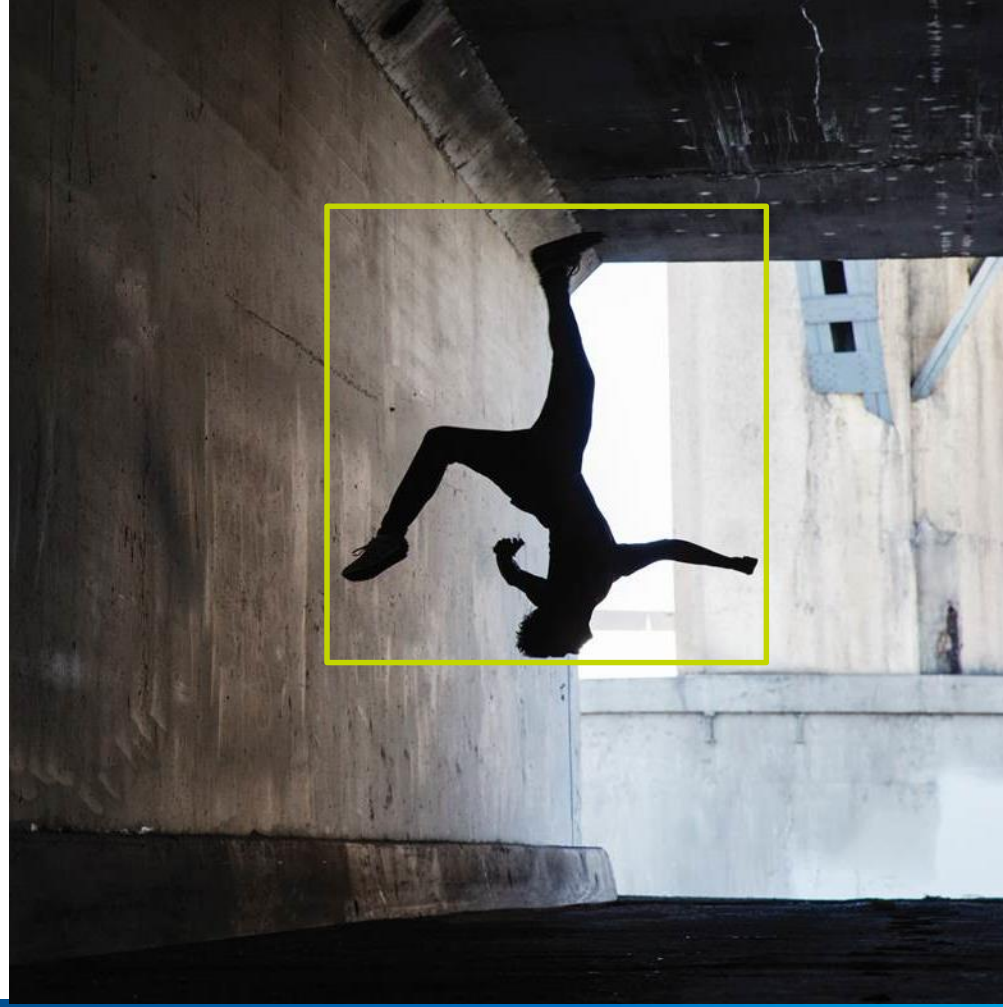

Activity 1: Setup Intel® Distribution of openVINO™ toolkit (openVINO)

Setup openVINO:

```
source /opt/intel/openvino/bin/setupvars.sh
```

Run Classification demo using SqueezeNet:

```
cd /opt/intel/openvino/deployment_tools/demo/
```

```
sudo ./demo_squeezenet_download_convert_run.sh
```

Run Security Barrier Camera Demo:

```
sudo ./demo_security_barrier_camera.sh
```

Activity 2: Model Optimizer

Download VGG_16

```
    cd /opt/intel/openvino/\
deployment_tools/tools/model_downloader/

    ./downloader.py --name vgg16 -o ~/lab5
```

Convert and Optimize model:

```
    cd /opt/intel/openvino/deployment_tools/model_optimizer/

    python36 mo.py --input_model ~/lab5/\
classification/vgg/16/caffe/vgg16.caffemodel \--output_dir ~/lab5/IR/ -
-reverse_input_channels -b 1
```

Activity 3: Classification

Run classification with VGG16:

```
    cd ~/inference_engine_samples_build/intel64/Release

   ./classification_sample \
-i ~/lab5/input_files/cat.jpg \
-m ~/lab5/IR/vgg16.xml
```

Activity 4: Object detection

Download model

```
cd /opt/intel/openvino/\
deployment_tools/tools/model_downloader/
```

```
./downloader.py --name ssd_mobilenet_v2_coco -o ~/lab5/
```

Convert and Optimize TensorFlow* model:

```
cd /opt/intel/openvino/deployment_tools/model_optimizer/
```

```
./mo_tf.py --input_model ~/lab5/object_detection/common/\
ssd_mobilenet_v2_coco/tf/ssd_mobilenet_v2_coco.frozen.pb \
--output_dir ~/lab5/IR/ --reverse_input_channels \
--tensorflow_use_custom_operations_config "extensions/front/tf/ssd_support.json" \
--tensorflow_object_detection_api_pipeline_config ~/lab5/tf_model/pipeline.config
```

Activity 4.1: Run video

Object detection with openVINO:

```
cd ~/inference_engine_samples_build/intel64/Release/


./object_detection_demo_ssd_async \
-i ~/lab5/input_files/denver.mp4 \
-m ~/lab5/IR/ssd_mobilenet_v2_coco.frozen.xml -t 0.6
```

Activity 5: VNNI

Download models:

```
cd /opt/intel/\

openvino/deployment_tools/tools/model_downloader/
```

Download FP32:

```
./downloader.py --name vehicle-\
attributes-recognition-barrier-0039 -o ~/lab6/
```

Download INT8:

```
./downloader.py --name vehicle-\
attributes-recognition-barrier-0039-int8 -o ~/lab6/
```

Activity 5.1: VNNI

Benchmak:

```
cd ~/inference_engine_samples_build/intel64/Release/
```

Download FP32:

```
./benchmark_app -m ~/lab6/Security/object_attributes/vehicle/\
resnet10_update_1/dldt/\
vehicle-attributes-recognition-barrier-0039.xml -i /opt/\
intel/openvino/deployment_tools/demo/car.png
```

Download INT8:

```
./benchmark_app -m ~/lab6/Security/object_attributes/vehicle/\
resnet10_update_1/dldt/\
vehicle-attributes-recognition-barrier-0039-int8.xml -i /opt/\
intel/openvino/deployment_tools/demo/car.png
```


intel
Software