



INTEL OPENVINO VIDEO INFERENCE ON THE EDGE

Igor Freitas



NOTICES AND DISCLAIMERS

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at www.intel.com.

Performance results are based on testing as of Aug. 20, 2017 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

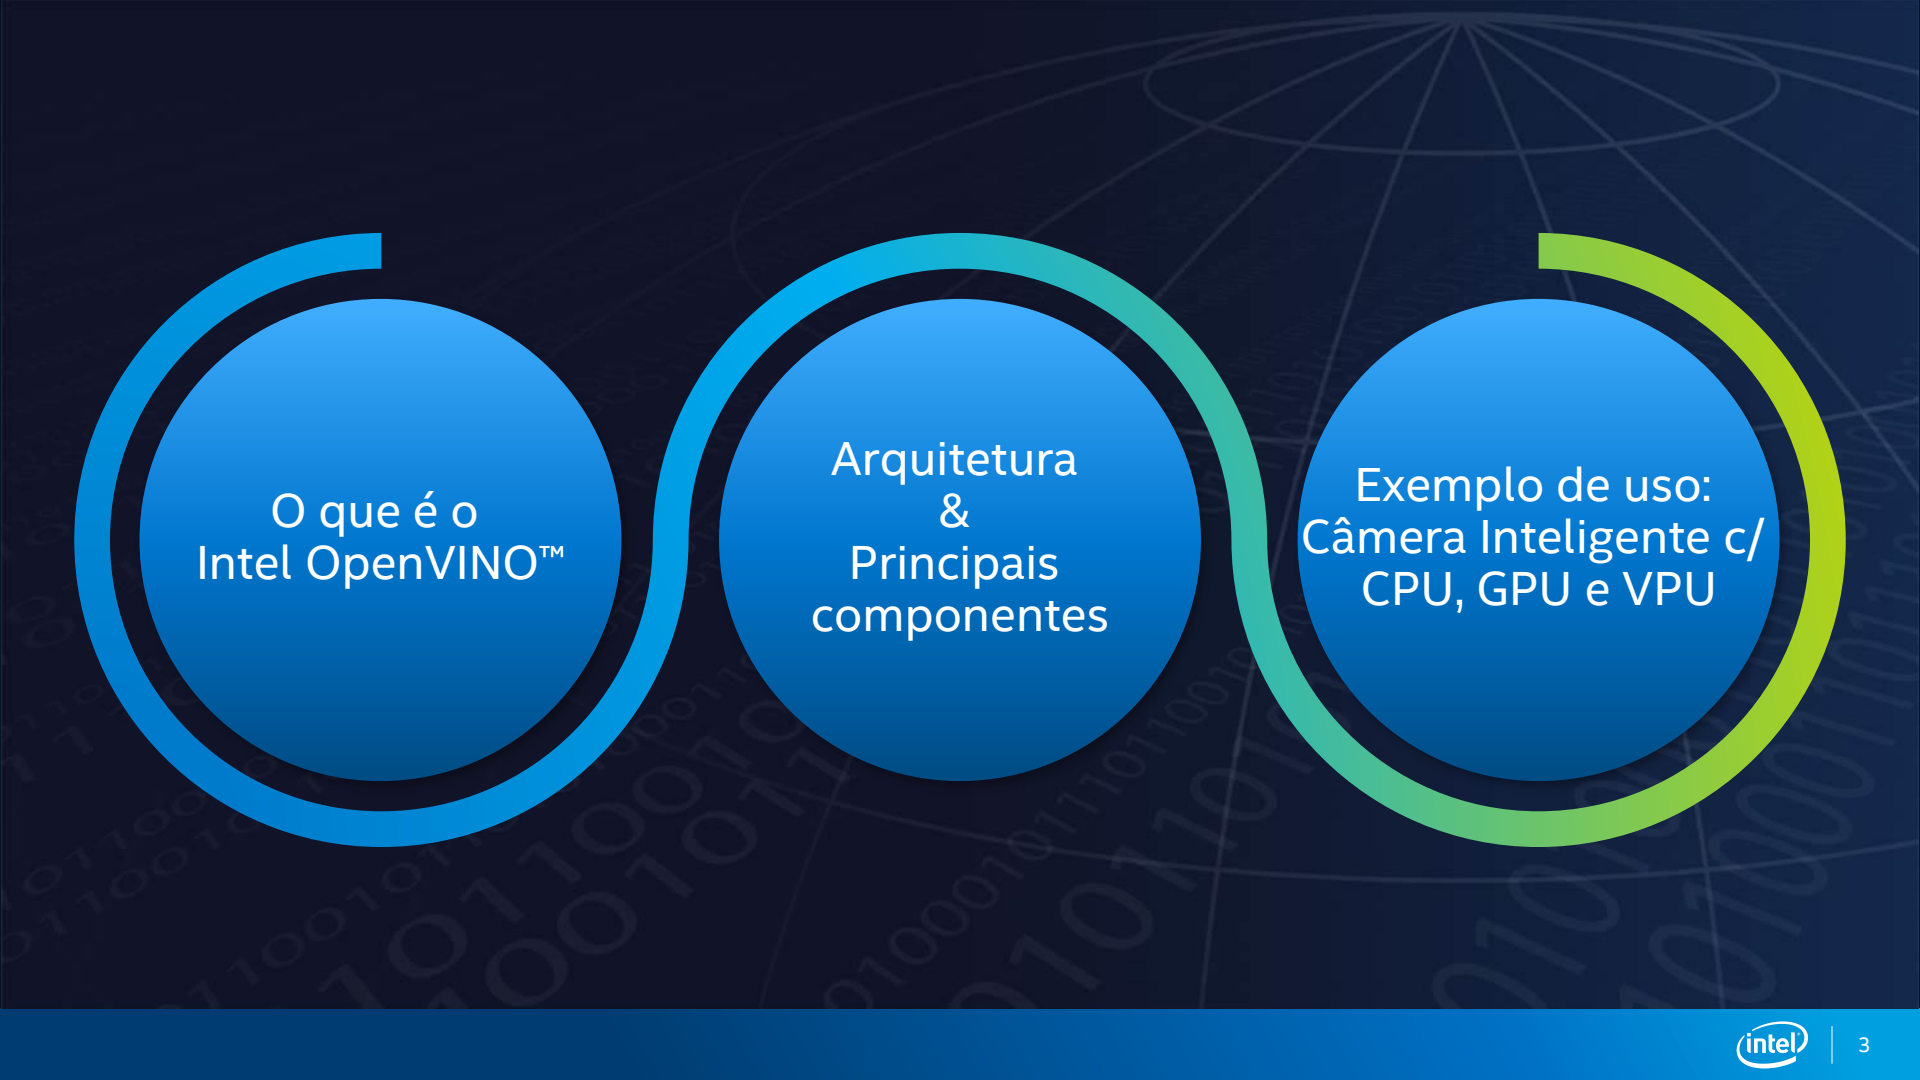
Any forecasts of goods and services needed for Intel's operations are provided for discussion purposes only. Intel will have no liability to make any purchase in connection with forecasts published in this document.

ARDUINO 101 and the ARDUINO infinity logo are trademarks or registered trademarks of Arduino, LLC.

Altera, Arria, the Arria logo, Intel, the Intel logo, Intel Atom, Intel Core, Intel Nervana, Intel Saffron, Iris, Movidius, OpenVINO, Stratix and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Copyright 2019 Intel Corporation.



O que é o
Intel OpenVINO™

Arquitetura
&
Principais
componentes

Exemplo de uso:
Câmera Inteligente c/
CPU, GPU e VPU

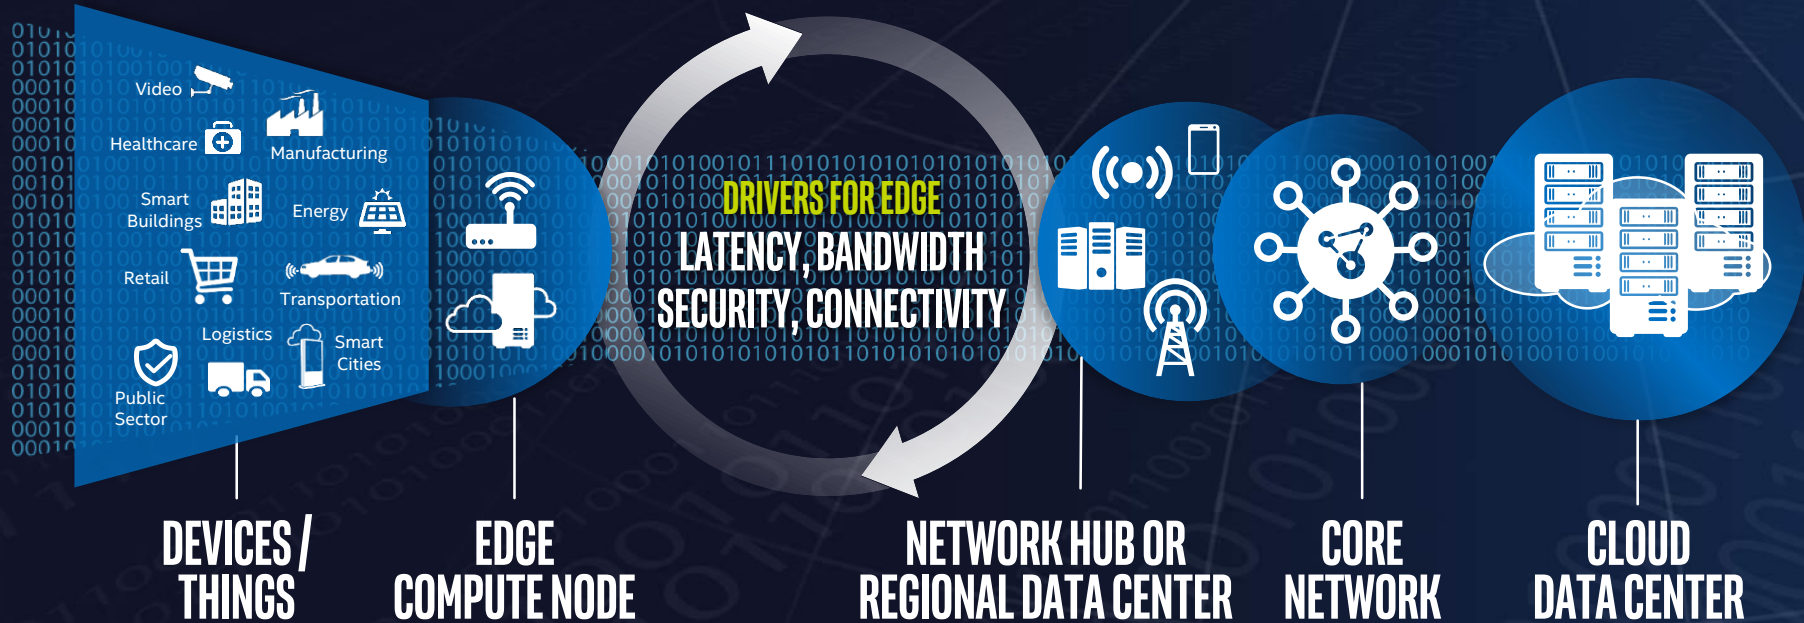


O que é o
Intel OpenVINO™

Arquitetura
&
Principais
componentes

Exemplo de uso:
Câmera Inteligente c/
CPU, GPU e VPU

THE EMERGING NEED FOR EDGE COMPUTE



INTEL® DISTRIBUTION OF OPENVINO™ TOOLKIT

WRITE ONCE, DEPLOY EVERYWHERE

DEEP LEARNING

Caffe TensorFlow ONNX mxnet KALDI

- Model Optimizer
- Inference Engine
- Supports >100 Public Models, incl. 30+ Pretrained Models
- CV Algorithms

COMPUTER VISION

OpenCV* OpenCL™ OpenVX™

- CV Library (Kernel & Graphic APIs)
- Optimized media encode/decode functions


AGNOSTIC, COMPLEMENTARY TO MAJOR FRAMEWORKS

CROSS-PLATFORM FLEXIBILITY

HIGH PERFORMANCE, HIGH EFFICIENCY



12,000+ Developers



Over 20 Customer Products Launched based on Intel® Distribution of OpenVINO™ toolkit



Breadth of vision product portfolio

STRONG ADOPTION + RAPIDLY EXPANDING CAPABILITY

O que é o
Intel OpenVINO™

Arquitetura
&
Principais
Componentes

Exemplo de uso:
Câmera Inteligente c/
CPU, GPU e VPU

INTEL® AI TOOLS

PORTFOLIO OF SOFTWARE TOOLS TO EXPEDITE AND ENRICH AI DEVELOPMENT

TOOLKITS



Application
Developers

DEEP LEARNING DEPLOYMENT

OpenVINO™†

Open Visual Inference & Neural Network Optimization toolkit for inference deployment on CPU/GPU/FPGA/VPU using TensorFlow*, for all Intel® Movidius™ VPUs using TensorFlow & Caffe* & MXNet*

Intel® Movidius™ SDK

Optimized inference deployment

DEEP LEARNING

COMING SOON!

Intel® Deep Learning Studio†

Open-source tool to compress deep learning development cycle

LIBRARIES



Data
Scientists

MACHINE LEARNING LIBRARIES

Python

- [Scikit-learn](#)
- [Pandas](#)
- [NumPy](#)

R

- [Cart](#)
- [Random Forest](#)
- [e1071](#)

Distributed

- [MLlib \(on Spark\)](#)
- [Mahout](#)

DEEP LEARNING FRAMEWORKS

Now optimized for CPU



[TensorFlow](#)



[MXNet](#)



[Caffe](#)



[BigDL* \(Spark\)](#)

Optimizations in progress



[Caffe2](#)



[PyTorch](#)



[CNTK](#)



[PaddlePaddle](#)

COMING SOON!

FOUNDATION



Library
Developers

ANALYTICS, MACHINE & DEEP LEARNING PRIMITIVES

Python*

Intel distribution optimized for machine learning

DAAL

Intel® Data Analytics Acceleration Library (incl machine learning)

MKL-DNN

Open-source deep neural network functions for CPU / integrated graphics

c1DNN

DEEP LEARNING GRAPH COMPILER

Intel® nGraph™ Compiler (Alpha)

Open-sourced compiler for deep learning model computations optimized for multiple devices from multiple frameworks

† Formerly the Intel® Computer Vision SDK

*Other names and brands may be claimed as the property of others.

Developer personas show above represent the primary user base for each row, but are not mutually-exclusive.

All products, computer systems, dates, and figures are preliminary based on current expectations, and are subject to change without notice.

INTEL® VISION PRODUCTS

WRITE ONCE - DEPLOY ACROSS INTEL ARCHITECTURE - LEVERAGE COMMON ALGORITHMS

OpenVINO™



Intel® CPUs
(Atom®, Core™, Xeon®)



Intel® CPUs
w/ Integrated Graphics



Intel® Movidius™ VPUs
& Intel® FPGAs



Future Accelerators
(Keem Bay, etc.)

1. Intel® Distribution of OpenVINO™ toolkit: Computer vision & deep learning inference tool with common API
2. Portfolio of hardware for computer vision & deep learning inference, device to cloud
3. Ecosystem to cover the breadth of IoT vision systems



INTEL® VISION ACCELERATOR DESIGN PRODUCTS

Add to existing Intel® architectures for accelerated DL inference capabilities

What's Inside Intel® Distribution of OpenVINO™ toolkit

Intel® Deep Learning Deployment Toolkit

Model Optimizer
Convert & Optimize



Inference Engine
Optimized Inference

30+ Pre-trained
Models

Computer Vision
Algorithms

Samples

IR = Intermediate Representation file



Traditional Computer Vision

Optimized Libraries & Code Samples

OpenCV*

OpenVX*

Samples

For Intel® CPU & GPU/Intel® Processor Graphics

Tools & Libraries

Increase Media/Video/Graphics Performance

Intel® Media SDK
Open Source version

OpenCL™
Drivers & Runtimes

For GPU/Intel® Processor Graphics

Optimize Intel® FPGA (Linux* only)

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

Bitstreams

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



Intel® Vision Accelerator
Design Products &
AI in Production/
Developer Kits




An open source version is available at 01.org/openvintoolkit (some deep learning functions support Intel CPU/GPU only).

Model Optimizer

- **What it is:** A python based tool to import trained models and convert them to Intermediate representation.
- **Why important:** Optimizes for performance/space with conservative topology transformations; biggest boost is from conversion to data types matching hardware.



IR = Intermediate Representation format

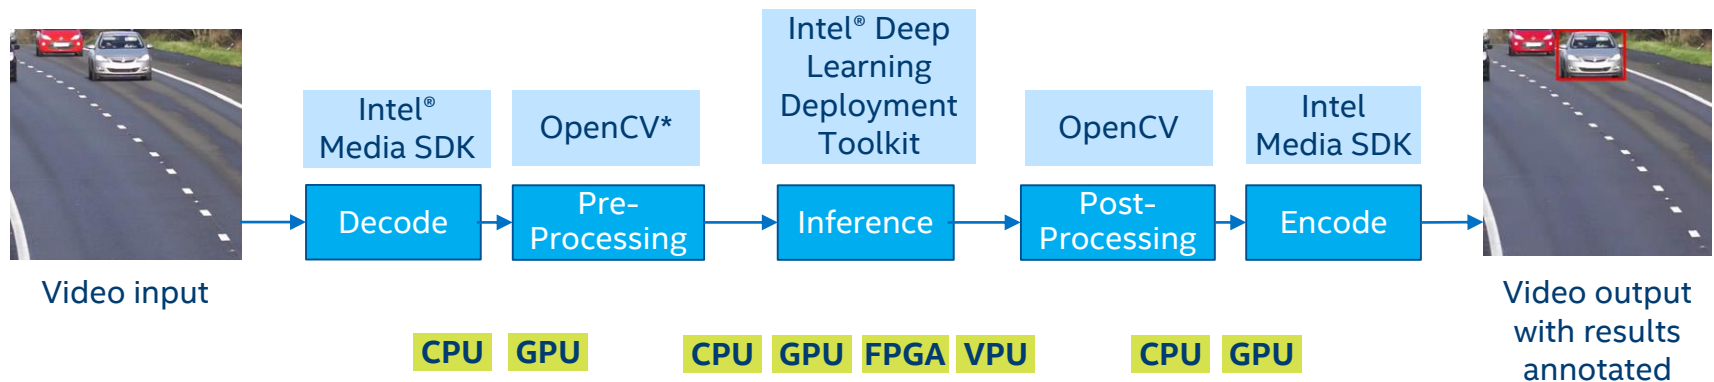


Inference Engine

- **What it is:** High-level inference API
- **Why important:** Interface is implemented as dynamically loaded plugins for each hardware type. Delivers best performance for each type without requiring users to implement and maintain multiple code pathways.

GPU = Intel CPU with integrated graphics processing unit/Intel® Processor Graphics

End-to-End Vision Workflow



O que é o
Intel OpenVINO™

Arquitetura
&
Principais
componentes

Exemplo de uso:
Câmera Inteligente c/
CPU, GPU e VPU

Classification




CAT

Classification + Localization



CAT

Object Detection



CAT, DOG, DUCK

Instance Segmentation



CAT, DOG, DUCK

Single object

Multiple objects

Comparison between image classification, object detection and instance segmentation.

Intel® AI Tools

Intel OpenVINO™ - Exemplo de uso

```
igor@igor-dell: ~
```

```
[setupvars.sh] OpenVINO environment initialized
```

```
(base) igor@igor-dell:~$ source /opt/intel/computer_vision_sdk/bin/setupvars.sh
```

```
igor@igor-dell: ~/tdc
```

```
(base) igor@igor-dell:~/tdc$ ls -all
```

```
total 32
```

```
drwxrwxr-x  2 igor igor 4096 Mar  1 09:41 .
drwxr-xr-x 45 igor igor 4096 Mar  1 09:38 ..
lrwxrwxrwx  1 igor igor  52 Mar  1 09:41 demo -> /opt/intel/computer_vision_sdk/deployment_tools/demo
lrwxrwxrwx  1 igor igor  61 Mar  1 09:41 documentation -> /opt/intel/computer_vision_sdk/deployment_tools/documentation
lrwxrwxrwx  1 igor igor  67 Mar  1 09:41 extension_generator -> /opt/intel/computer_vision_sdk/deployment_tools/extension_generator
lrwxrwxrwx  1 igor igor  64 Mar  1 09:40 inference_engine -> /opt/intel/computer_vision_sdk/deployment_tools/inference_engine
lrwxrwxrwx  1 igor igor  60 Mar  1 09:40 intel_models -> /opt/intel/computer_vision_sdk/deployment_tools/intel_models
lrwxrwxrwx  1 igor igor  78 Mar  1 09:39 model_downloader -> /opt/intel/computer_vision_sdk/deployment_tools/model_downloader/downloader.py
lrwxrwxrwx  1 igor igor  69 Mar  1 09:39 model_optimizer -> /opt/intel/computer_vision_sdk/deployment_tools/model_optimizer/mo.py
(base) igor@igor-dell:~/tdc$
```

Intel® AI Tools

Intel OpenVINO™ – Model Downloader

```
(openvino) igor@igor-dell:~/tdc$ cat computer_vision_sdk/deployment_tools/model_downloader/list_topologies.yml | grep name
(openvino) igor@igor-dell:~/tdc$ cat computer_vision_sdk/deployment_tools/model_downloader/list_topologies.yml | grep name
- name: "densenet-121"
- name: "densenet-161"
- name: "densenet-169"
- name: "densenet-201"
- name: "squeezeNet1.0"
- name: "squeezeNet1.1"
- name: "mtcnn-p"
- name: "mtcnn-r"
- name: "mtcnn-o"
- name: "mobilenet-ssd"
- name: "vgg19"
- name: "vgg16"
- name: "ssd512"
- name: "ssd300"
- name: "inception-resnet-v2"
- name: "dilation"
- name: "googlenet-v1"
- name: "googlenet-v2"
- name: "googlenet-v4"
- name: "alexnet"
- name: "ssd_mobilenet_v2_coco"
- name: "resnet-50"
- name: "resnet-101"
- name: "resnet-152"
- name: "googlenet-v3"
- name: "se-inception"
- name: "se-resnet-101"
- name: "se-resnet-152"
- name: "se-resnet-50"
- name: "se-resnext-50"
- name: "se-resnext-101"
- name: "Sphereface"
- name: "license-plate-recognition-barrier-0007"
- name: "age-gender-recognition-retail-0013"
- name: "age-gender-recognition-retail-0013-fp16"
- name: "emotions-recognition-retail-0003"
- name: "emotions-recognition-retail-0003-fp16"
- name: "face-detection-adas-0001"
- name: "face-detection-adas-0001-fp16"
- name: "face-detection-retail-0004"
- name: "face-detection-retail-0004-fp16"
- name: "face-person-detection-retail-0002"
- name: "face-person-detection-retail-0002-fp16"
- name: "face-reidentification-retail-0071"
- name: "face-reidentification-retail-0071-fp16"
- name: "facial-landmarks-35-adas-0001"
- name: "facial-landmarks-35-adas-0001-fp16"
- name: "head-pose-estimation-adas-0001"
- name: "head-pose-estimation-adas-0001-fp16"
- name: "human-pose-estimation-0001"
- name: "human-pose-estimation-0001-fp16"
- name: "landmarks-regression-retail-0009"
- name: "landmarks-regression-retail-0009-fp16"
- name: "license-plate-recognition-barrier-0001"
- name: "license-plate-recognition-barrier-0001-fp16"
- name: "pedestrian-and-vehicle-detector-adas-0001"
- name: "pedestrian-and-vehicle-detector-adas-0001-fp16"
- name: "pedestrian-detection-adas-0002"
- name: "pedestrian-detection-adas-0002-fp16"
- name: "person-attributes-recognition-crossroad-0031"
```


Intel® AI Tools

Intel OpenVINO™ – Model Downloader

```
igor@igor-dell: ~/tdc
(openvino) igor@igor-dell:~/tdc$ python model_downloader -h
usage: model_downloader [-h] [-c CONFIG] [--name NAME] [--all] [--print_all]
                        [-o OUTPUT_DIR]
```

optional arguments:

```
-h, --help            show this help message and exit
-c CONFIG, --config CONFIG
                        path to YML configuration file
--name NAME           names of topologies for downloading with comma
                        separation
--all                 download all topologies from the configuration file
--print_all           print all available topologies
-o OUTPUT_DIR, --output_dir OUTPUT_DIR
                        path where to save topologies
```

list_topologies.yml - default configuration file

```
(openvino) igor@igor-dell:~/tdc$
```

```
python model_downloader.py --name ssd512 --o .
```

```
(openvino) igor@igor-dell:~/tdc$ python model_downloader.py --name ssd512 -o .
```

```
#####|| Start downloading models ||#####
```

```
#####|| Start downloading weights ||#####
```

```
#####|| Start downloading topologies in tarballs ||#####
```

```
...100%, 98624 KB, 15849 KB/s, 6 seconds passed ===== ssd512.tar.gz =====> ./object_detection/common/ssd/512/caffe/ssd512.tar.gz
```

Intel® AI Tools

Intel OpenVINO™ – Model Optimizer

```
(openvino) igor@igor-dell:~/tdc$ python model_optimizer.py -h
usage: model_optimizer.py [-h] [--framework {tf,caffe,mxnet,kaldi,onnx}]
                          [--input_model INPUT_MODEL]
                          [--model_name MODEL_NAME] [--output_dir OUTPUT_DIR]
                          [--input_shape INPUT_SHAPE] [--scale SCALE]
                          [--reverse_input_channels]
                          [--log_level {CRITICAL,ERROR,WARN,WARNING,INFO,DEBUG,NOTSET}]
                          [--input INPUT] [--output OUTPUT]
                          [--mean_values MEAN_VALUES]
                          [--scale_values SCALE_VALUES]
                          [--data_type {FP16,FP32,half,float}]
                          [--disable_fusing] [--disable_resnet_optimization]
                          [--finegrain_fusing FINEGRAIN_FUSING]
                          [--disable_gfusing] [--move_to_preprocess]
                          [--extensions EXTENSIONS] [--batch BATCH]
                          [--version] [--silent]
                          [--freeze_placeholder_with_value FREEZE_PLACEHOLDER_WITH_VALUE]
                          [--generate_deprecated_IR_V2]
                          [--input_model_is_text]
                          [--input_checkpoint INPUT_CHECKPOINT]
                          [--input_meta_graph INPUT_META_GRAPH]
                          [--saved_model_dir SAVED_MODEL_DIR]
                          [--saved_model_tags SAVED_MODEL_TAGS]
                          [--offload_unsupported_operations_to_tf]
                          [--tensorflow_subgraph_patterns TENSORFLOW_SUBGRAPH_PATTERNS]
                          [--tensorflow_operation_patterns TENSORFLOW_OPERATION_PATTERNS]
                          [--tensorflow_custom_operations_config_update TENSORFLOW_CUSTOM_OPERATIONS_CONFIG_UPDATE]
                          [--tensorflow_use_custom_operations_config TENSORFLOW_USE_CUSTOM_OPERATIONS_CONFIG]
                          [--tensorflow_object_detection_api_pipeline_config TENSORFLOW_OBJECT_DETECTION_API_PIPELINE_CONFIG]
                          [--tensorboard_logdir TENSORBOARD_LOGDIR]
                          [--tensorflow_custom_layer_libraries TENSORFLOW_CUSTOM_LAYER_LIBRARIES]
                          [--disable_nhwc_to_nchw] [--input_proto INPUT_PROTO]
                          [-k K] [--mean_file MEAN_FILE]
                          [--mean_file_offsets MEAN_FILE_OFFSETS]
                          [--disable_omitting_optional]
                          [--enable_flattening_nested_params]
                          [--input_symbol INPUT_SYMBOL]
                          [--nd_prefix_name ND_PREFIX_NAME]
                          [--pretrained_model_name PRETRAINED_MODEL_NAME]
                          [--save_params_from_nd] [--legacy_mxnet_model]
                          [--counts COUNTS] [--remove_output_softmax]
```

python model_optimizer.py --h

Intel® AI Tools

Intel OpenVINO™ – Model Optimizer

```
(openvino) igor@igor-dell:~/tdc$ python model_optimizer.py --input_model /home/igor/tdc/object_detection/common/ssd/300/caffe/ssd300.caffemodel
Model Optimizer arguments:
Common parameters:
- Path to the Input Model:      /home/igor/tdc/object_detection/common/ssd/300/caffe/ssd300.caffemodel
- Path for generated IR:       /home/igor/tdc/.
- IR output name:              ssd300
- Log level:                    ERROR
- Batch:                        Not specified, inherited from the model
- Input layers:                 Not specified, inherited from the model
- Output layers:               Not specified, inherited from the model
- Input shapes:                 Not specified, inherited from the model
- Mean values:                  Not specified
- Scale values:                 Not specified
- Scale factor:                 Not specified
- Precision of IR:              FP32
- Enable fusing:                True
- Enable grouped convolutions fusing:  True
- Move mean values to preprocess section:  False
- Reverse input channels:       False
Caffe specific parameters:
- Enable resnet optimization:    True
- Path to the Input prototxt:    /home/igor/tdc/object_detection/common/ssd/300/caffe/ssd300.prototxt
- Path to CustomLayersMapping.xml:  Default
- Path to a mean file:           Not specified
- Offsets for a mean file:       Not specified
Model Optimizer version:        1.5.12.49d067a0

[ SUCCESS ] Generated IR model.
[ SUCCESS ] XML file: /home/igor/tdc/./ssd300.xml
[ SUCCESS ] BIN file: /home/igor/tdc/./ssd300.bin
[ SUCCESS ] Total execution time: 7.47 seconds.
(openvino) igor@igor-dell:~/tdc$
```

```
python model_optimizer.py --input_model /path/ssd300.caffemodel
```

Intel® AI Tools

Intel OpenVINO™ – Model Optimizer – Movidius ou GPU – 16FP

```
(pv) igor@igor-dell:~/tdc$ python model_optimizer.py --input_model /home/igor/tdc/object_detection/common/ssd/300/caffe/ssd300.caffemodel --data_type FP16 -o ssd300_openvino/FP16/
^[[3-Model Optimizer arguments:
Common parameters:
- Path to the Input Model:      /home/igor/tdc/object_detection/common/ssd/300/caffe/ssd300.caffemodel
- Path for generated IR:       /home/igor/tdc/ssd300_openvino/FP16/
- IR output name:              ssd300
- Log level:                    ERROR
- Batch:                        Not specified, inherited from the model
- Input layers:                 Not specified, inherited from the model
- Output layers:               Not specified, inherited from the model
- Input shapes:                 Not specified, inherited from the model
- Mean values:                  Not specified
- Scale values:                 Not specified
- Scale factor:                 Not specified
- Precision of IR:              FP16
- Enable fusing:                True
- Enable grouped convolutions fusing: True
- Move mean values to preprocess section: False
- Reverse input channels:      False
Caffe specific parameters:
- Enable resnet optimization:   True
- Path to the Input prototxt:   /home/igor/tdc/object_detection/common/ssd/300/caffe/ssd300.prototxt
- Path to CustomLayersMapping.xml: Default
- Path to a mean file:          Not specified
- Offsets for a mean file:      Not specified
Model Optimizer version:        1.5.12.49d067a0
[ WARNING ]
Detected not satisfied dependencies:
  protobuf: installed: 3.5.2, required: 3.6.1

Please install required versions of components or use install_prerequisites script
/opt/intel/computer_vision_sdk_2018.5.445/deployment_tools/model_optimizer/install_prerequisites/install_prerequisites_caffe.sh
Note that install_prerequisites scripts may install additional components.

[ SUCCESS ] Generated IR model.
[ SUCCESS ] XML file: /home/igor/tdc/ssd300_openvino/FP16/ssd300.xml
[ SUCCESS ] BIN file: /home/igor/tdc/ssd300_openvino/FP16/ssd300.bin
[ SUCCESS ] Total execution time: 7.41 seconds.
```

```
python model_optimizer.py --input_model /path/ssd300.caffemodel --
ata_type FP16 --o /path
```

Intel® AI Tools

Intel OpenVINO™ – Inference Engine – Async Demo

```
(openvino) igor@igor-dell:~/tdc$ python computer_vision_sdk/inference_engine/samples/python_samples/object_detection_demo_ssd_async.py
/home/igor/anaconda3/envs/openvino/lib/python3.6/importlib/_bootstrap.py:219: RuntimeWarning: compiletime version 3.5 of module 'openvino.inference_engine.ie_api' does not match runtime version 3.6
  return f(*args, **kwargs)
usage: object_detection_demo_ssd_async.py [-h] -m MODEL -i INPUT
                                         [-l CPU_EXTENSION] [-pp PLUGIN_DIR]
                                         [-d DEVICE] [--labels LABELS]
                                         [-pt PROB_THRESHOLD]
object_detection_demo_ssd_async.py: error: the following arguments are required: -m/--model, -i/--input
(openvino) igor@igor-dell:~/tdc$
```

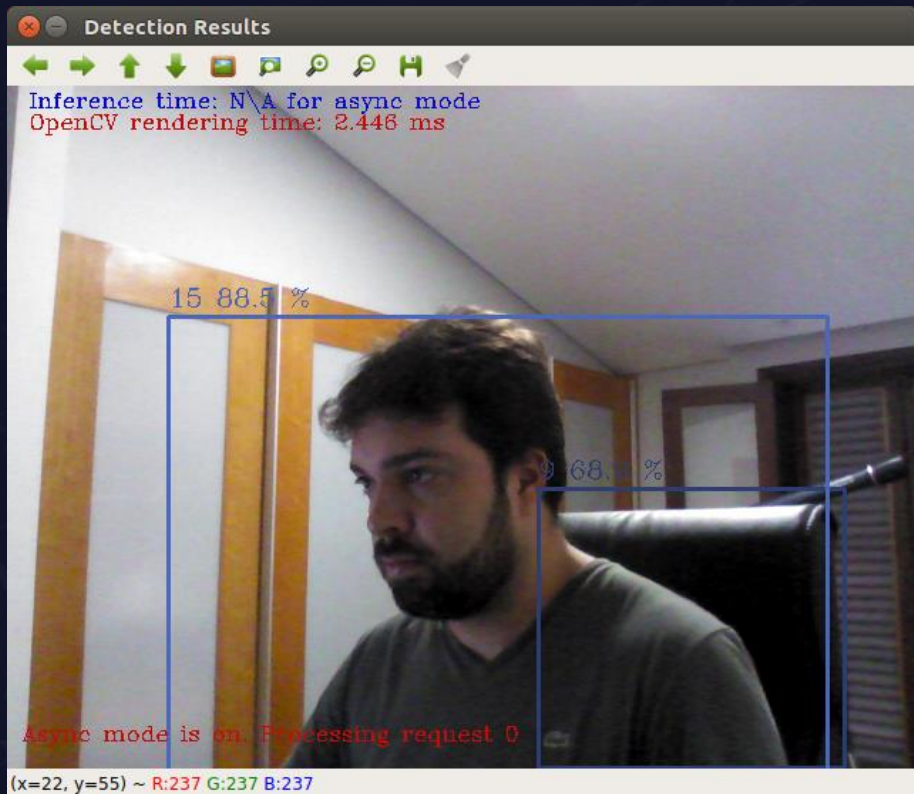
```
(pv) igor@igor-dell:~/tdc$ python computer_vision_sdk/deployment_tools/inference_engine/samples/python_samples/object_detection_demo_ssd_async.py -m /home/igor/tdc/ssd300_openvino/ssd300.xml -i cam -d MYRIAD
/home/igor/anaconda3/envs/pv/lib/python3.6/importlib/_bootstrap.py:219: RuntimeWarning: compiletime version 3.5 of module 'openvino.inference_engine.ie_api' does not match runtime version 3.6
  return f(*args, **kwargs)
[ INFO ] Initializing plugin for MYRIAD device...
[ INFO ] Reading IR...
[ INFO ] Loading IR to the plugin...
Traceback (most recent call last):
  File "computer_vision_sdk/deployment_tools/inference_engine/samples/python_samples/object_detection_demo_ssd_async.py", line 182, in <module>
    sys.exit(main() or 0)
  File "computer_vision_sdk/deployment_tools/inference_engine/samples/python_samples/object_detection_demo_ssd_async.py", line 77, in main
    exec_net = plugin.load(network=net, num_requests=2)
  File "ie_api.pyx", line 389, in openvino.inference_engine.ie_api.IEPlugin.load
  File "ie_api.pyx", line 400, in openvino.inference_engine.ie_api.IEPlugin.load
RuntimeError: [VPU] Unsupported network precision : FP32
(pv) igor@igor-dell:~/tdc$
```

➔ VPU ou GPU requer FP16 bits

```
python object_detection_demo.py -m /path-FP16/ssd300.xml -i cam -d MYRIAD
```


Intel® AI Tools

Intel OpenVINO™ – Inference Engine – Async Demo



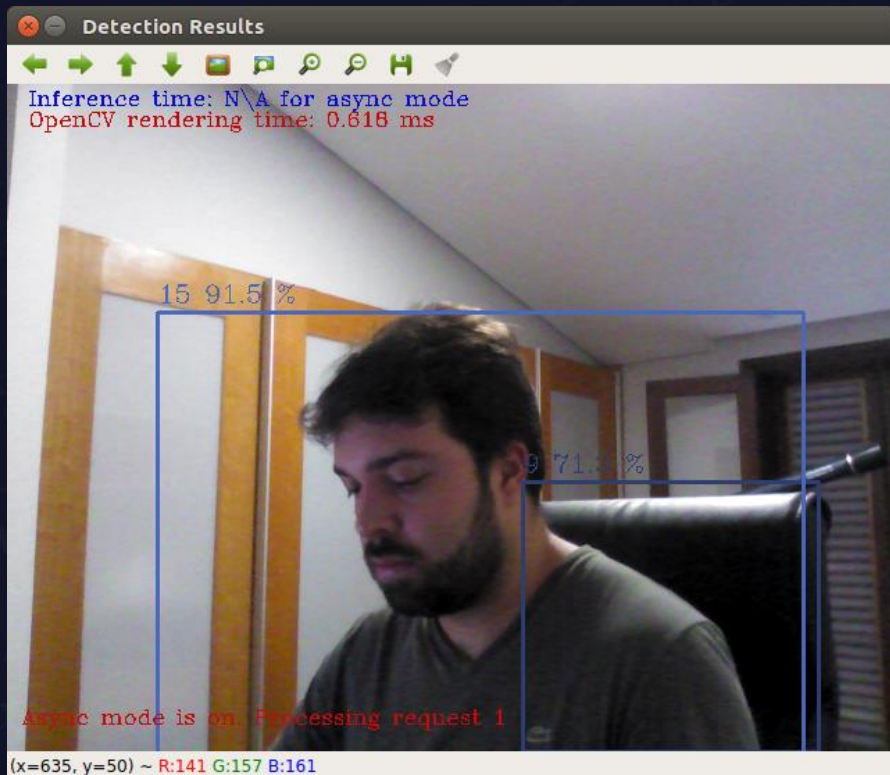
Intel® Movidius™ Myriad™ VPU

Render time: 2.4ms (416 FPS)



Intel® AI Tools

Intel OpenVINO™ – Inference Engine – Async Demo



Intel Integrated GPU Running
Modelo
Render time: 0.6ms (1.666 FPS)

Call to Action, Resources

Download ▶

[Free Intel® Distribution of
OpenVINO™ toolkit](#)


Get started quickly with:

- [Developer resources](#)
- [Intel® Tech.Decoded online webinars, tool
how-tos & quick tips](#)
- [Hands-on developer workshops](#)

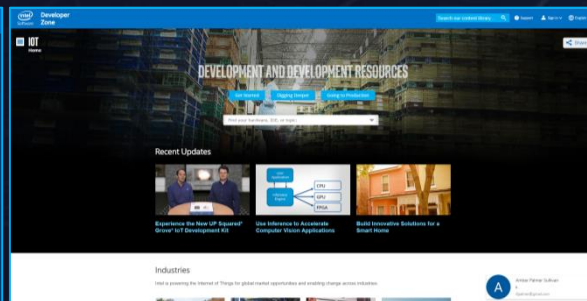
Support

- 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.



GET CONNECTED



[WWW.INTEL.COM/IOT](http://www.intel.com/iot)

[AI.INTEL.COM](http://ai.intel.com)

[SOFTWARE.INTEL.COM/IOT/ARDUINO-CREATE](http://software.intel.com/iot/arduino-create)

[SOFTWARE.INTEL.COM/SYSTEM-STUDIO](http://software.intel.com/system-studio)

[SOFTWARE.INTEL.COM/SDKS](http://software.intel.com/sdks)

[WWW.INTEL.COM/VISIONPRODUCTS](http://www.intel.com/visionproducts)

[WWW.INTEL.COM/MARKETREADYSURVEY](http://www.intel.com/marketreadysurvey)

[WWW.INTEL.COM/RFPKITS](http://www.intel.com/rfpkits)

What's New in Intel® Distribution of OpenVINO™ toolkit 2018 R5

- **Extends neural network support** to include LSTM (long short-term memory) from ONNX*, TensorFlow*& MXNet* frameworks, & 3D convolutional-based networks in preview mode (CPU-only) for **non-vision use cases**.
- **Introduces Neural Network Builder API** (preview), providing flexibility to create a graph from simple API calls and directly deploy via the Inference Engine.
- **Improves Performance** - Delivers significant CPU performance boost on multicore systems through **new parallelization techniques via streams**. Optimizes performance on Intel® Xeon®, Core™ & Atom processors through **INT8-based primitives** for Intel® Advanced Vector Extensions (Intel® AVX-512), Intel® AVX2 & SSE4.2.
- **Supports Raspberry Pi* hardware** as a host for the Intel® Neural Compute Stick 2 (preview). Offload your deep learning workloads to this low-cost, low-power USB.
- **Adds 3 new optimized pretrained models** (for a total of 30+): Text detection of indoor/outdoor scenes, and 2 single-image super resolution networks that enhance image resolution by a factor of 3 or 4.

See product site & release notes for more details about 2018 R4.

