


## PRODUCT BRIEF

Computer Vision and Deep Learning  
Intel® Distribution of OpenVINO™ Toolkit



# Accelerate Computer Vision and Deep Learning from Edge to Cloud

## Fast-Track Development of High-Performance Computer Vision and Deep Learning Inference into Vision Applications



Intel® Distribution of OpenVINO™ toolkit (open visual inference and neural network optimization) is free software that helps developers and data scientists speed up computer vision workloads, streamline deep learning inference and deployments, and enable easy, heterogeneous execution across Intel® platforms from edge to cloud. It helps to:

- **Increase deep learning workload performance up to 19x<sup>1</sup>** with computer vision accelerators from Intel. Enable heterogeneous processing and asynchronous execution across multiple types of Intel® processors—CPUs, CPUs with integrated graphics, Intel® FPGAs, and Intel® Movidius™ vision processing units (VPUs).
- **Unleash convolutional neural network (CNN)-based deep learning inference** using a common API. Streamline deep learning inference and deployment using standard or custom layers without the overhead of frameworks. Includes 20+ pre-trained models, ready-to-use Computer-Vision Algorithms (CVA), and support for 100+ open source and custom models.
- **Speed development** using optimized OpenCV\* and OpenVX\* functions. Get started quickly with 15+ code samples. Add custom kernels into workloads (i.e., video/image processing, computer vision routines, feature extraction/tracking, OpenCL™).

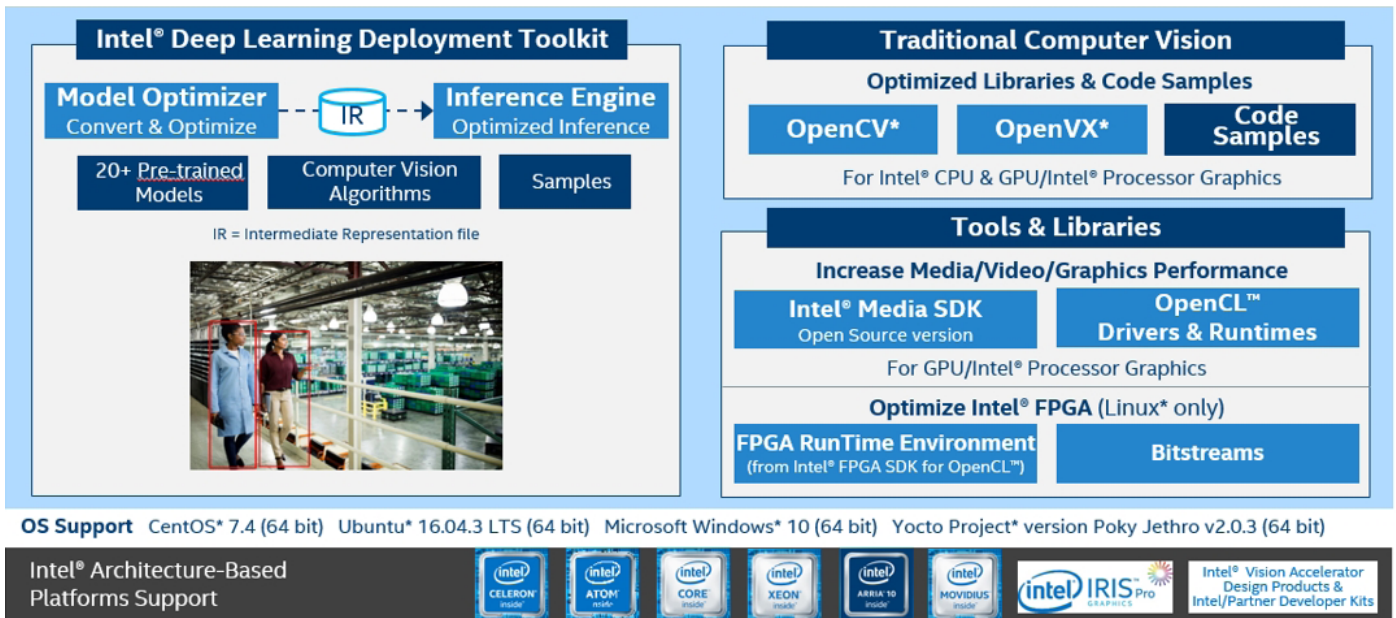
Demand for intelligent vision solutions is increasing from edge to cloud—including digital surveillance, retail, industrial/manufacturing, smart cities and homes, office automation, autonomous driving, and more. Deep learning revenue is expected to grow from \$655 million in 2016 to \$35 billion by 2025<sup>2</sup>.

Top systems require multiple fast, scalable, reliable platforms and devices to support multi-channel streaming, real-time software-based analytics deep learning, and more. Intel's robust hardware and software products give OEMs/ODMs, system integrators, ISVs, and solution providers the tools they need to speed design, development, and deployment of high-performance computer vision solutions—enabling rich solutions for AI everywhere.

### What's Inside

- **Intel® Deep Learning Deployment Toolkit** with model optimizer and inference engine
- **Optimized computer vision libraries** for OpenCV and OpenVX for CPUs with integrated graphics (Intel® Processor Graphics)
- **Components to increase performance of Intel Processor Graphics for Linux\***, including the Intel® Media SDK open source version and OpenCL graphics drivers and runtimes
- **FPGA Runtime Environment** (from the Intel® FPGA SDK for OpenCL™) and bitstreams for Linux FPGA

**Free Download >**  
[software.intel.com/opencvino-toolkit](https://software.intel.com/opencvino-toolkit)



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

Figure 1. Includes Intel® Deep Learning Deployment Toolkit with model optimizer, inference engine, and other component tools and libraries.

## Key Components

### Intel® Deep Learning Deployment Toolkit

#### Model Optimizer

**What it does:** Imports trained models from various frameworks (Caffe\*, Tensorflow\*, MxNet\*, ONNX\*, Kaldi\*) and converts them to a unified intermediate representation file. It also optimizes topologies through node merging, horizontal fusion, eliminating batch normalization, and quantization. It also supports graph freeze and graph summarize along with dynamic input freezing.

**Why it's important:** The Model Optimizer typically provides the biggest performance boost by conversion to data types that match hardware types (FP32 to FP16).

#### Includes samples for:

- Standard and pipelined image classification
- Image segmentation
- Object detection
- Object detection for Single Shot Multibox Detector (SSD)
- Neural style transfer
- Security barrier
- Interactive face detection
- People counter
- Multi-channel face detection

#### Inference Engine

**What it is:** A simple and unified API for inference across Intel® architecture. It allows execution of layers across hardware targets (CPU, CPU with integrated graphics, FPGA, and VPU) through heterogeneous support.

**Why it's important:** The inference engine's interface is implemented as dynamically loaded plugins for each hardware type. It delivers the best performance for each type without making you implement and maintain multiple code pathways.

#### Pre-trained models:

- Age/gender
- Face detection (standard and enhanced)
- Face re-identification
- Head position
- Human detection (eye-level and high-angle detection)
- Detect people, vehicles, and bikes
- License plate detection (small and front-facing)
- Vehicle metadata
- Pedestrian and vehicle detection
- Retail environment
- Person attributes recognition crossroad
- Person and action detection for smart classroom
- Emotion recognition
- Identify someone from different videos
- Advanced roadside identification
- Vehicle attributes
- Landmarks regression
- Crossroad object detection
- Semantic segmentation
- Facial landmarks
- Human pose estimation
- Single image super resolution

#### Optimized Libraries and Functions for Traditional Computer Vision

Get faster performance across heterogeneous Intel® platforms: CPU, CPU with integrated graphics, FPGA, Intel® Movidius™ VPUs:

- **OpenCV:** A well-established, open source, and optimized computer vision library. It includes a wide variety of algorithms and functions available including the Intel® Photography Vision Library, a collection of highly optimized and state of the art computer vision functions.
- **OpenVX:** An optimized, graph-based approach for computer vision functions that is targeted at real-time, low power applications. More than 15 samples included.

## Why You Need It

### Boost Computer Vision Performance

- Provides access to Intel's computer vision accelerators (CPU, CPU with integrated graphics, FPGAs, and Intel Movidius VPUs) to accelerate performance.
- Supports heterogeneous execution across multiple types of Intel® processors
- Accelerates algorithms

### Streamline Deep Learning Inference and Deployment

- Enables CNN-based deep learning inference on edge devices.
- Includes Deep Learning Deployment Toolkit with a model optimizer and an inference engine, which provides optimized inference across many hardware types.
- Allows you to use standard or custom layers without the overhead of frameworks.
- Includes more than 20 pre-trained models, six optimized and ready to deploy Computer Vision algorithms, and support for 100+ open source and custom models.

### Speed Development for Vision Solutions, Save Time with a Heterogeneous Approach

- Supports heterogeneous execution across multiple types of Intel processors.
- Improves efficiency and saves development time.
- Use the SDK's library of functions, pre-optimized kernels, and optimized calls for OpenCV and OpenVX. Write code once and deploy it to multiple types of Intel processor-based devices, easily customizing it for your models. Delivers broad capabilities to all developers—no domain expertise required.


## Accelerate Computer Vision and Deep Learning Inference from Edge to Cloud

Download the free Intel Distribution of OpenVINO toolkit to speed up computer vision workloads and streamline deep learning deployments with easy, heterogeneous execution across multiple types of Intel hardware.

Optimize more capabilities by using the Intel Distribution of OpenVINO toolkit with other tools:

- **Intel® SDK for OpenCL™ Applications** for Intel® CPUs and CPUs with integrated graphics workload balancing
- **Intel® System Studio** to optimize system bring-up and IOT device application performance

## Success Stories



### Deep Learning for Healthcare Imaging

Intel and GE brought the power of AI to clinical diagnostic scanning and other healthcare workflows.

[Read more >](#)



### Performance Boost for a Deep Learning Algorithm

GeoVision sped up its facial recognition solution using Intel® System Studio and the Intel Distribution of OpenVINO toolkit.

[Read more >](#)

[See more success stories >](#)

## Get Started Now

- [Download the free Intel® Distribution of OpenVINO™ toolkit >](#)
- [Learn more about Intel Distribution of OpenVINO toolkit >](#)
- [Explore Intel® Deep Learning Deployment Toolkit >](#)
- [Visit Tech.Decoded, Intel's developer knowledge hub >](#)

## Technical Specifications

<b>CPU</b>	<p><b>Supported Hardware</b></p> <ul style="list-style-type: none"> <li>6th-8th generation Intel® Core™ processors with Iris® Pro graphics and Intel® HD Graphics</li> <li>Intel® Pentium® processor N4200/5, N3350/5, N3450/5 with Intel HD Graphics</li> </ul> <p><b>Supported Operating Systems</b></p> <ul style="list-style-type: none"> <li>Microsoft Windows® 10 (64 bit)</li> <li>Ubuntu* 16.04.3 LTS (64 bit)</li> <li>CentOS* 7.4 (64 bit)</li> <li>Yocto Project* MR3 (64 bit)</li> </ul> <p><b>Supported OpenVINO Toolkit Components</b></p> <ul style="list-style-type: none"> <li>Intel® Deep Learning Deployment Toolkit</li> <li>OpenCV</li> <li>OpenVX*</li> </ul>
<b>CPU with Integrated Graphics (Intel® Processor Graphics/GPU)</b>	<p><b>Supported Hardware</b></p> <ul style="list-style-type: none"> <li>6th-8th generation Intel® Core™ processors with Iris® Pro graphics and Intel® HD Graphics</li> <li>Intel® Pentium® processor N4200/5, N3350/5, N3450/5 with Intel HD Graphics</li> </ul> <p><b>Supported Operating Systems</b></p> <ul style="list-style-type: none"> <li>Windows 10 (64 bit)</li> <li>Ubuntu 16.04.3 LTS (64 bit)</li> <li>CentOS 7.4 (64 bit)</li> <li>Yocto Project MR3 (64 bit)</li> </ul> <p><b>Supported OpenVINO Toolkit Components</b></p> <ul style="list-style-type: none"> <li>Intel® Deep Learning Deployment Toolkit</li> <li>OpenCV</li> <li>OpenVX</li> </ul>
<b>Intel® FPGAs</b>	<p><b>Supported Hardware</b></p> <ul style="list-style-type: none"> <li>Intel® Arria® 10 GX FPGA Development Kit</li> <li>Intel® Programmable Acceleration Card with Intel® Arria® 10 GX FPGA</li> </ul> <p><b>Supported Operating Systems</b></p> <ul style="list-style-type: none"> <li>Ubuntu 16.04.3 LTS (64 bit)</li> <li>CentOS 7.4 (64 bit)</li> </ul> <p><b>Supported OpenVINO Toolkit Components</b></p> <ul style="list-style-type: none"> <li>Intel® Deep Learning Deployment Toolkit</li> </ul>
<b>Intel® Movidius™ Vision Processing Unit (VPU)</b>	<p><b>Supported Hardware</b></p> <ul style="list-style-type: none"> <li>Intel® Movidius™ Neural Compute Stick</li> <li>Intel® Neural Compute Stick 2</li> </ul> <p><b>Supported Operating Systems</b></p> <ul style="list-style-type: none"> <li>Ubuntu 16.04.3 TLS (64 bit)</li> <li>CentOS 7.4 (64 bit)</li> <li>Windows 10 (64 bit)</li> </ul> <p><b>Supported OpenVINO Toolkit Components</b></p> <ul style="list-style-type: none"> <li>Intel Deep Learning Deployment Toolkit</li> <li>OpenCV and OpenVX functions (needs to be run against the CPU or GPU)</li> </ul>
<b>Intel® Vision Accelerator Design</b>	<p><b>Supported Hardware:</b> Intel® Vision Accelerator Design with Intel® Arria10 FPGA (preview)</p> <p><b>Supported Operating Systems:</b> Ubuntu 16.04.3 TLS (64 bit)</p> <p><b>Supported OpenVINO Toolkit Components:</b> Intel® Deep Learning Deployment Toolkit</p>



<sup>1</sup> See image on p. 3 for details and configuration.

<sup>2</sup> Tractica, 2Q 2017

Performance results are based on testing as of 4/13/2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "re-served" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information. The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

Copyright © 2018 Intel Corporation. Intel, the Intel logo, Intel Core and Intel Xeon are trademarks of Intel Corporation in the U.S. and other countries.

\* Other names and brands may be claimed as the property of others. OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Kronos.