INTEL® INNOVATION DAY

intel®

THE PATH TO DEEPER INSIGHT

Advanced Analytics

Operational Analytics

Cognitive Analytics

Self–Learning
How Do I Proceed?

Prescriptive Analytics

Forecast
How Should I Proceed?

Predictive Analytics

Foresight
What Will Happen, When, and Why

Diagnostic Analytics

Insight
What Happened and Why?

Descriptive Analytics

Hindsight
What Happened?

AI
IS THE DRIVING FORCE
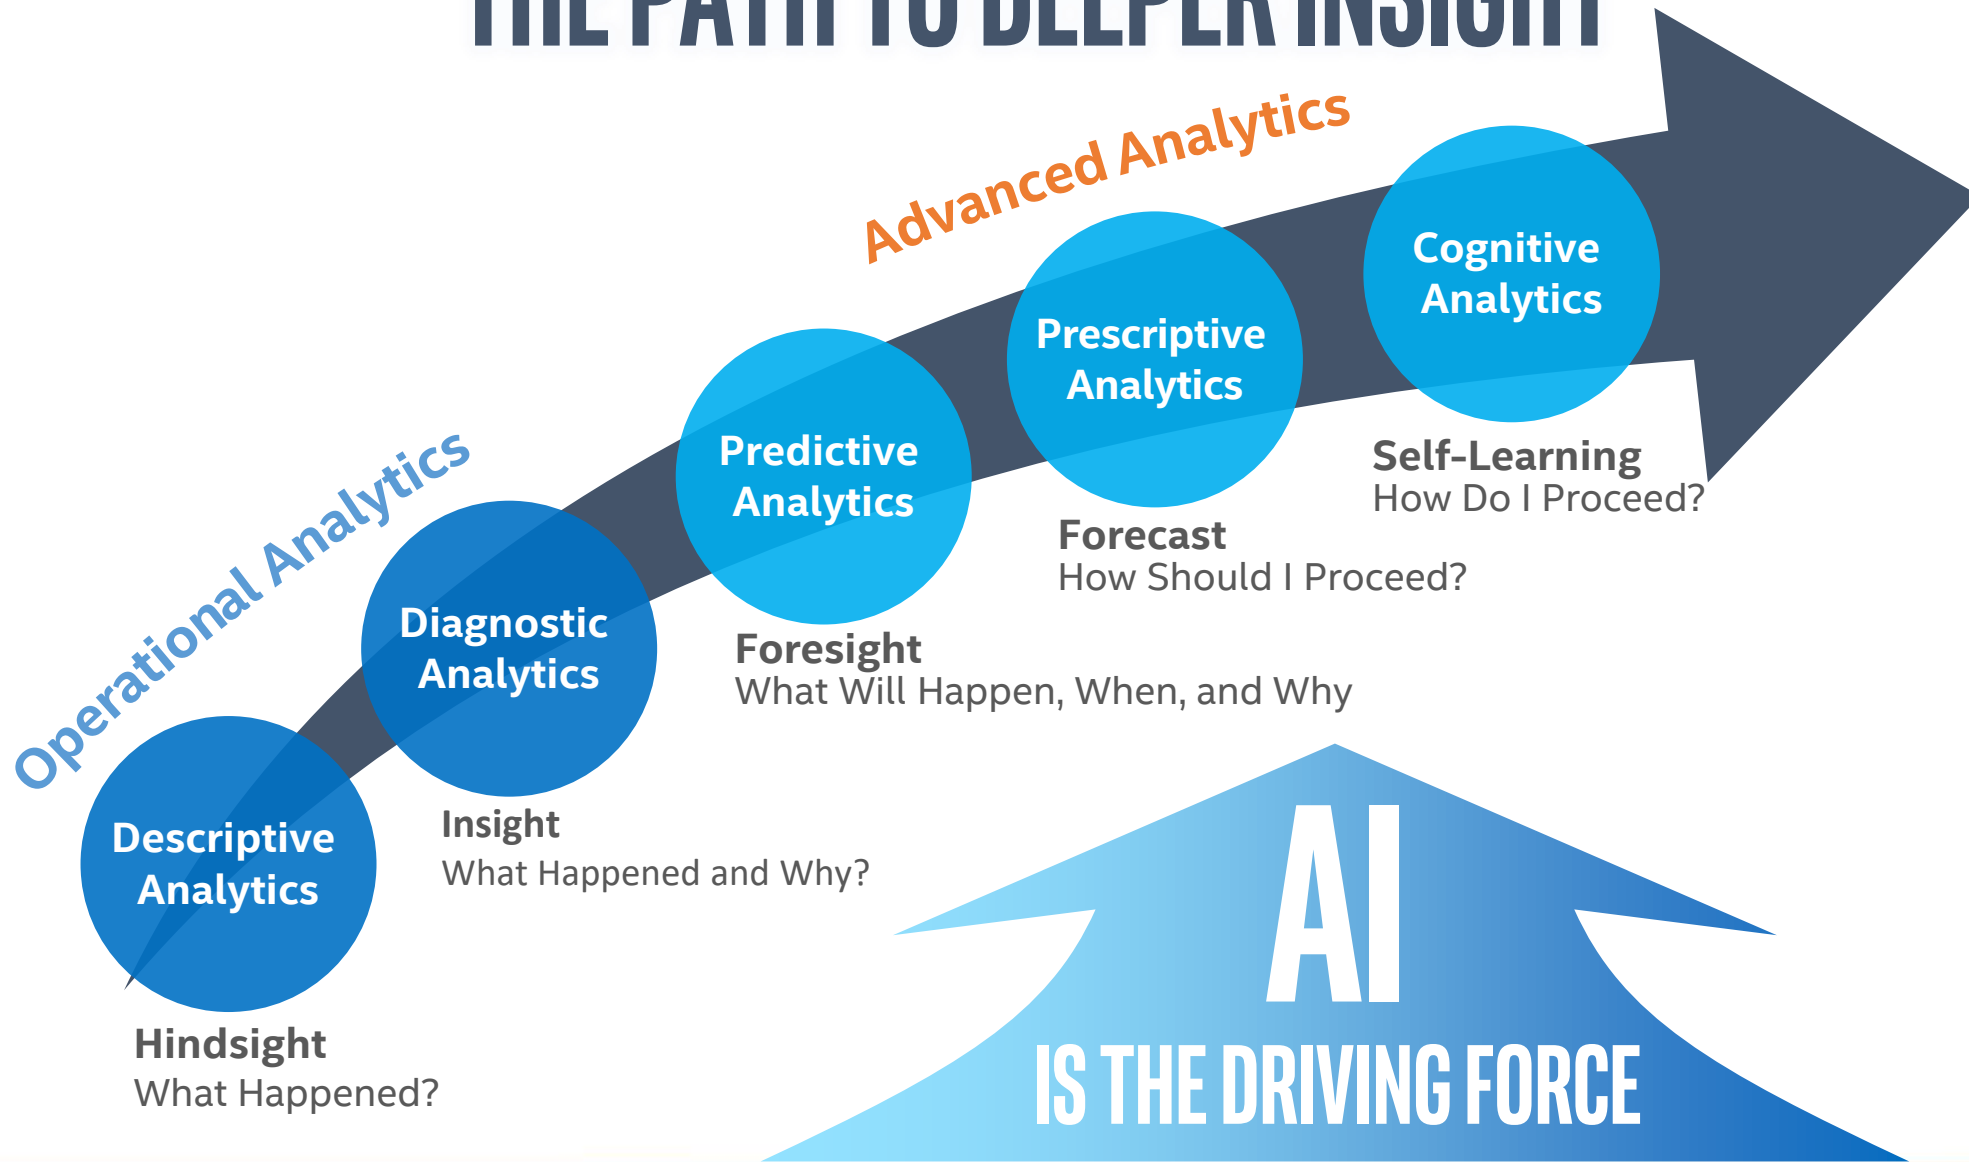
AI CLOSER LOOK

MACHINE LEARNING

Algorithms designed to deliver better insight with more data

Regression (Linear/Logistic)

Classification (Support Vector Machines/SVM, Naïve Bayes)

Clustering (Hierarchical, Bayesian, K-Means, DBSCAN)

Decision Trees (RandomForest)

Extrapolation (Hidden Markov Models/HMM)

More…
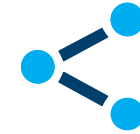
DEEP LEARNING

Neural networks used to infer meaning from large dense datasets

Image Recognition (Convolutional Neural Networks/CNN, Single-Shot Detector/SSD)

Speech Recognition (Recurrent Neural Network/RNN)

Natural Language Processing (Long-Short Term Memory/LSTM)

Data Generation (Generative Adversarial Networks/GAN)
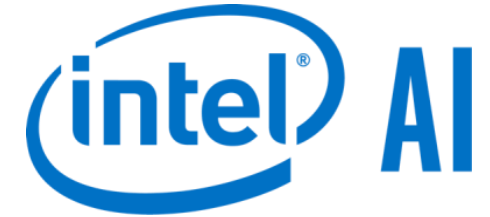
Recommender System (Multi-Layer Perceptron/MLP)

Time-Series Analysis (LSTM, RNN)

Reinforcement Learning (CNN, RNN)

More…

TOOLS
Software to accelerate development and deployment of real solutions

intel® AI

TOOLKITS
App Developers

DEEP LEARNING DEPLOYMENT

OpenVINO™ †
Open Visual Inference & Neural Network Optimization toolkit for inference deployment on CPU/GPU/FPGA for TF, Caffe* & MXNet*

Intel® Movidius™ SDK
Optimized inference deployment on Intel VPUs for TensorFlow* & Caffe*

DEEP LEARNING
COMING SOON!
Intel® Deep Learning Studio‡
Open-source tool to compress deep learning development cycle

LIBRARIES
Data Scientists

MACHINE LEARNING LIBS

Python
- Scikit-learn
- Pandas
- NumPy

R
- Cart
- RandomForest
- e1071

Distributed
- MlLib (on Spark)
- Mahout

DEEP LEARNING FRAMEWORKS
COMING SOON!

Now optimized for CPU

TensorFlow*   MXNet*   Caffe*   BigDL/Spark*

Optimizations in progress

Caffe2*   PyTorch*   PaddlePaddle*

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

DEEP LEARNING GRAPH COMPILER

Intel® nGraph™ Compiler (Alpha)
Open-sourced compiler for deep learning model computations optimized for multiple devices (CPU, GPU, NNP) from multiple frameworks (TF, MXNet, ONNX)
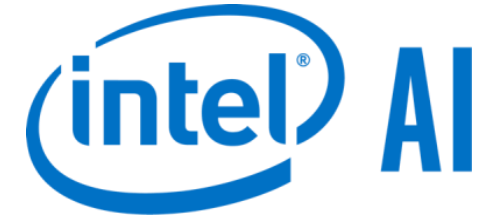
† Formerly the Intel® Computer Vision SDK
*Other names and brands may be claimed as the property of others.
Developer personas show above represent the primary user base for each row, but are not mutually-exclusive
All products, computer systems, dates, and figures are preliminary based on current expectations, and are subject to change without notice.

TOOLS

Software to accelerate development and deployment of real solutions

intel® AI

TOOLKITS
App Developers

DEEP LEARNING DEPLOYMENT		DEEP LEARNING
OpenVINO™[†]	**Intel® Movidius™ SDK**	**Intel® Deep Learning Studio‡**
Open Visual Inference & Neural Network Optimization toolkit for inference deployment on CPU/GPU/FPGA for TF, Caffe & MXNet**	*Optimized inference deployment on Intel VPUs for TensorFlow* & Caffe**	*Open-source tool to compress deep learning development cycle*

COMING SOON!

LIBRARIES
Data Scientists

MACHINE LEARNING LIBS			DEEP LEARNING FRAMEWORKS	
Python	**R**	**Distributed**	**Now optimized for CPU**	**Optimizations in progress**
•Scikit-learn	•Cart	•MlLib (on Spark)		
•Pandas	•RandomForest	•Mahout		
•NumPy	•e1071			

Now optimized for CPU: TensorFlow* MXNet* Caffe* BigDL/Spark*

Optimizations in progress: Caffe2* PyTorch* PaddlePaddle*

COMING SOON!

FOUNDATION
Library Developers

ANALYTICS, MACHINE & DEEP LEARNING PRIMITIVES			DEEP LEARNING GRAPH COMPILER
Python	**DAAL**	**MKL-DNN**	**Intel® nGraph™ Compiler** (Alpha)
Intel distribution optimized for machine learning	*Intel® Data Analytics Acceleration Library (incl machine learning)*	*Open-source deep neural network functions for CPU / integrated graphics*	*Open-sourced compiler for deep learning model computations optimized for multiple devices (CPU, GPU, NNP) from multiple frameworks (TF, MXNet, ONNX)*

† 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.

OPTIMIZED DEEP LEARNING FRAMEWORKS

Install an Intel-optimized framework and featured topology

More under optimization: Caffe2* PYTORCH* PaddlePaddle* n and more...

Get started today at ai.intel.com/framework-optimizations/

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

TOOLS

Software to accelerate development and deployment of real solutions

intel AI

TOOLKITS
App Developers

DEEP LEARNING DEPLOYMENT

OpenVINO™ †
Open Visual Inference & Neural Network Optimization toolkit for inference deployment on CPU/GPU/FPGA for TF, Caffe & MXNet**

Intel® Movidius™ SDK
Optimized inference deployment on Intel VPUs for TensorFlow & Caffe**

DEEP LEARNING
Intel® Deep Learning Studio‡
Open-source tool to compress deep learning development cycle

COMING SOON!

LIBRARIES
Data Scientists

MACHINE LEARNING LIBS
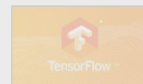
Python
- Scikit-learn
- Pandas
- NumPy

R
- Cart
- RandomForest
- e1071

Distributed
- MlLib (on Spark)
- Mahout

DEEP LEARNING FRAMEWORKS

Now optimized for CPU
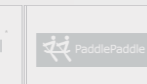
TensorFlow*   MXNet*   Caffe*   BigDL/Spark*

Optimizations in progress

Caffe2*   PyTorch*   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

DEEP LEARNING GRAPH COMPILER

Intel® nGraph™ Compiler (Alpha)
Open-sourced compiler for deep learning model computations optimized for multiple devices (CPU, GPU, NNP) from multiple frameworks (TF, MXNet, ONNX)

INTEL® INNOVATION DAY

3

INTEL® MKL-DNN

Intel's Open-Source Math Kernel Library for Deep Neural Networks
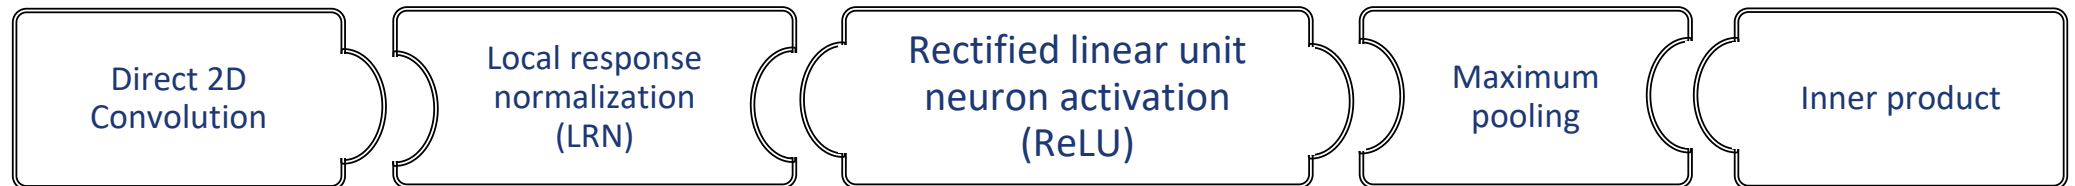
For developers of deep learning frameworks featuring optimized performance on Intel hardware

Distribution Details

- Open Source

- Apache 2.0 License

- Common DNN APIs across all Intel hardware.

- Rapid release cycles, iterated with the DL community, to best support industry framework integration.

- Highly vectorized & threaded for maximal performance, based on the popular Intel® MKL library.

github.com/01org/mkl-dnn

Examples:

| Direct 2D Convolution | Local response normalization (LRN) | Rectified linear unit neuron activation (ReLU) | Maximum pooling | Inner product |

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

INTEL® DATA ANALYTICS ACCELERATION LIBRARY (INTEL® DAAL)

High Performance Machine Learning and Data Analytics Library

Building blocks for all data analytics stages, including data preparation, data mining & machine learning

Pre-processing	→	Transformation	→	Analysis	→	Modeling	→	Validation	→	Decision Making

Open Source • Apache 2.0 License

Common Python, Java and C++ APIs across all Intel hardware

Optimized for large data sets including streaming and distributed processing

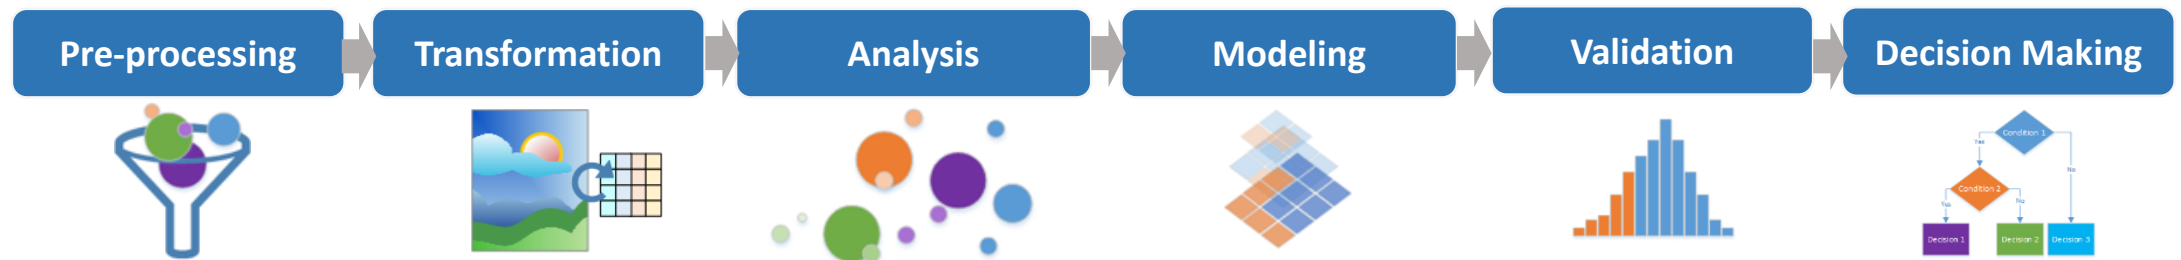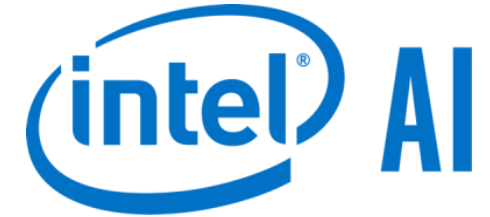Flexible interfaces to leading big data platforms including Spark and range of data formats (CSV, SQL, etc.)

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

INTEL® INNOVATION DAY

TOOLS

Software to accelerate development and deployment of real solutions

intel® AI

TOOLKITS
App Developers

DEEP LEARNING DEPLOYMENT		DEEP LEARNING **COMING SOON!**
OpenVINO™ † *Open Visual Inference & Neural Network Optimization toolkit for inference deployment on CPU/GPU/FPGA for TF, Caffe* & MXNet**	**Intel® Movidius™ SDK** *Optimized inference deployment on Intel VPUs for TensorFlow* & Caffe**	**Intel® Deep Learning Studio‡** *Open-source tool to compress deep learning development cycle*

LIBRARIES
Data Scientists

MACHINE LEARNING LIBS

Python	**R**	**Distributed**
• Scikit-learn • Pandas • NumPy	• Cart • RandomForest • e1071	• MlLib (on Spark) • Mahout

DEEP LEARNING FRAMEWORKS

Now optimized for CPU   **Optimizations in progress** **COMING SOON!**

TensorFlow*   MXNet*   Caffe*   BigDL/Spark*   Caffe2*   PyTorch*   PaddlePaddle*

FOUNDATION
Library Developers

ANALYTICS, MACHINE & DEEP LEARNING PRIMITIVES

Python	**DAAL**	**MKL-DNN**
Intel distribution optimized for machine learning	*Intel® Data Analytics Acceleration Library (incl machine learning)*	*Open-source deep neural network functions for CPU / integrated graphics*

DEEP LEARNING GRAPH COMPILER

Intel® nGraph™ Compiler (Alpha)

Open-sourced compiler for deep learning model computations optimized for multiple devices (CPU, GPU, NNP) from multiple frameworks (TF, MXNet, ONNX)
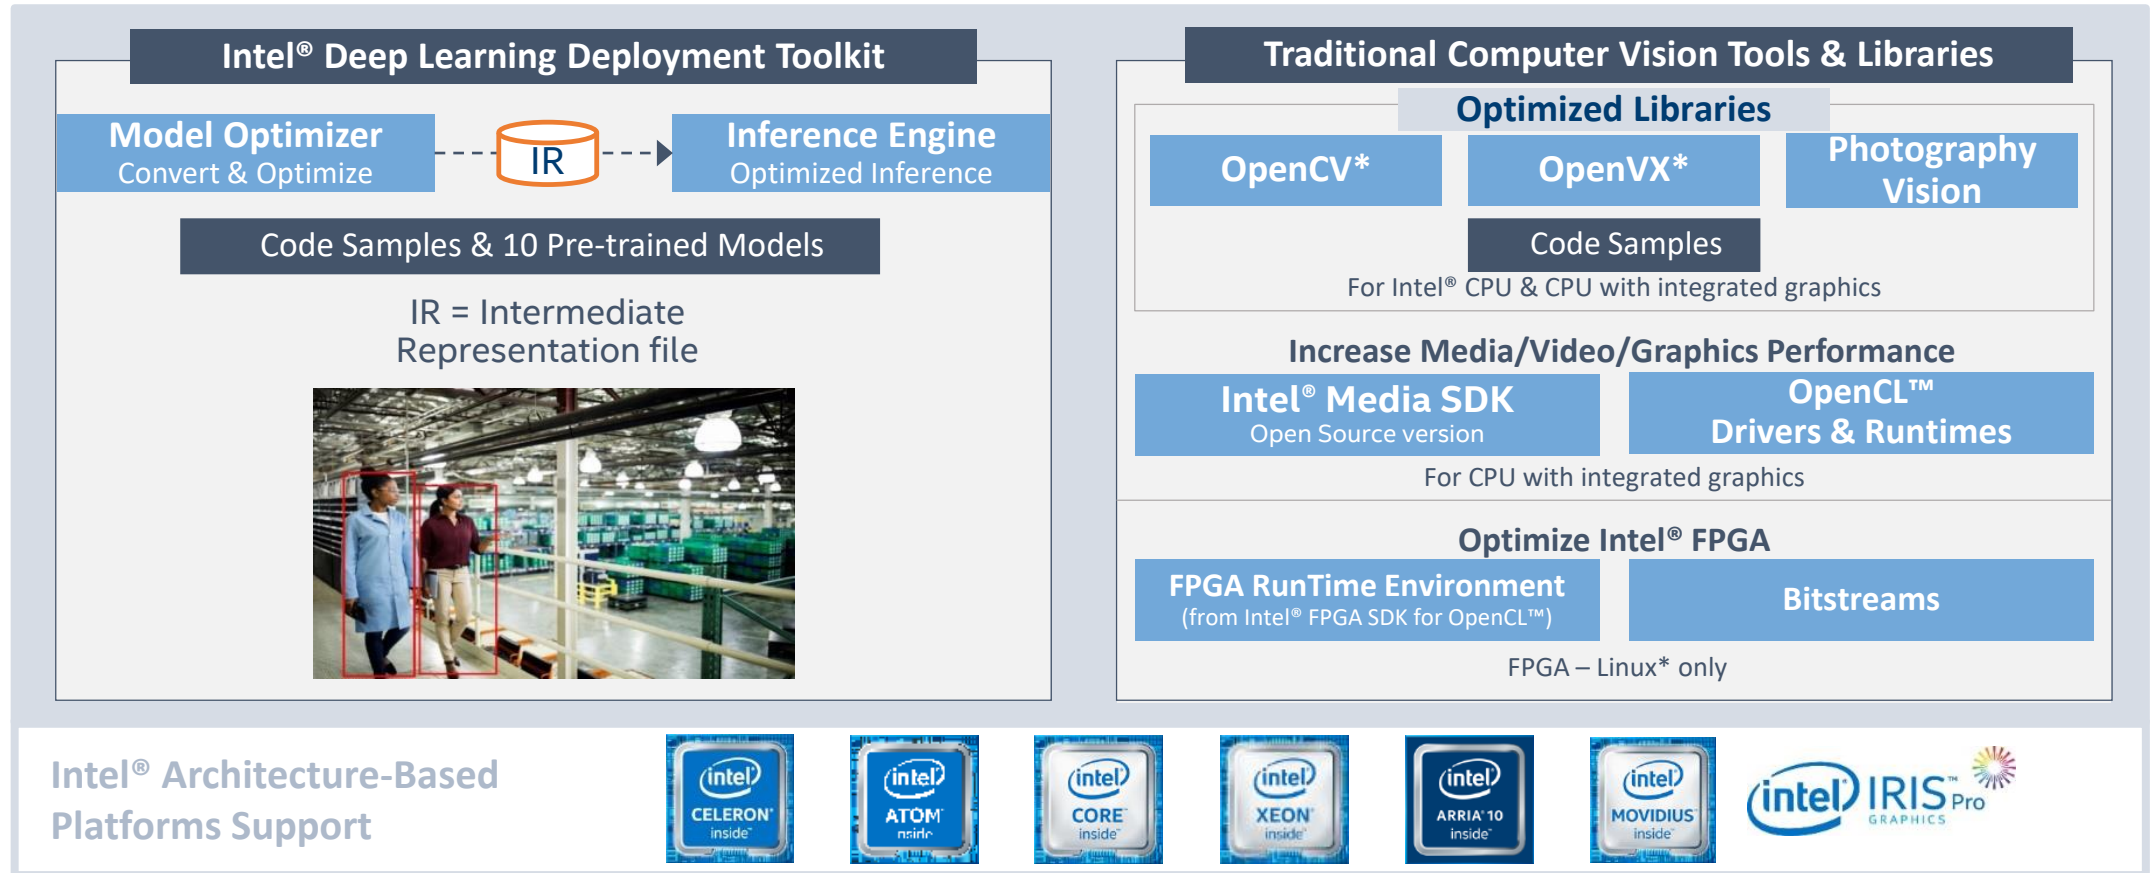
† Formerly the Intel® Computer Vision SDK
*Other names and brands may be claimed as the property of others.
Developer personas show above represent the primary user base for each row, but are not mutually-exclusive
All products, computer systems, dates, and figures are preliminary based on current expectations, and are subject to change without notice.

OPENVINO™ TOOLKIT

Cross-Platform Tool to Accelerate Computer Vision & Deep Learning Inference Performance

Intel® Deep Learning Deployment Toolkit

Model Optimizer
Convert & Optimize

IR

Inference Engine
Optimized Inference

Code Samples & 10 Pre-trained Models

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

Intel® Architecture-Based Platforms Support

intel CELERON inside

intel ATOM inside

intel CORE inside

intel XEON inside

intel ARRIA 10 inside

intel MOVIDIUS inside

intel IRIS Pro GRAPHICS

software.intel.com/openvino-toolkit

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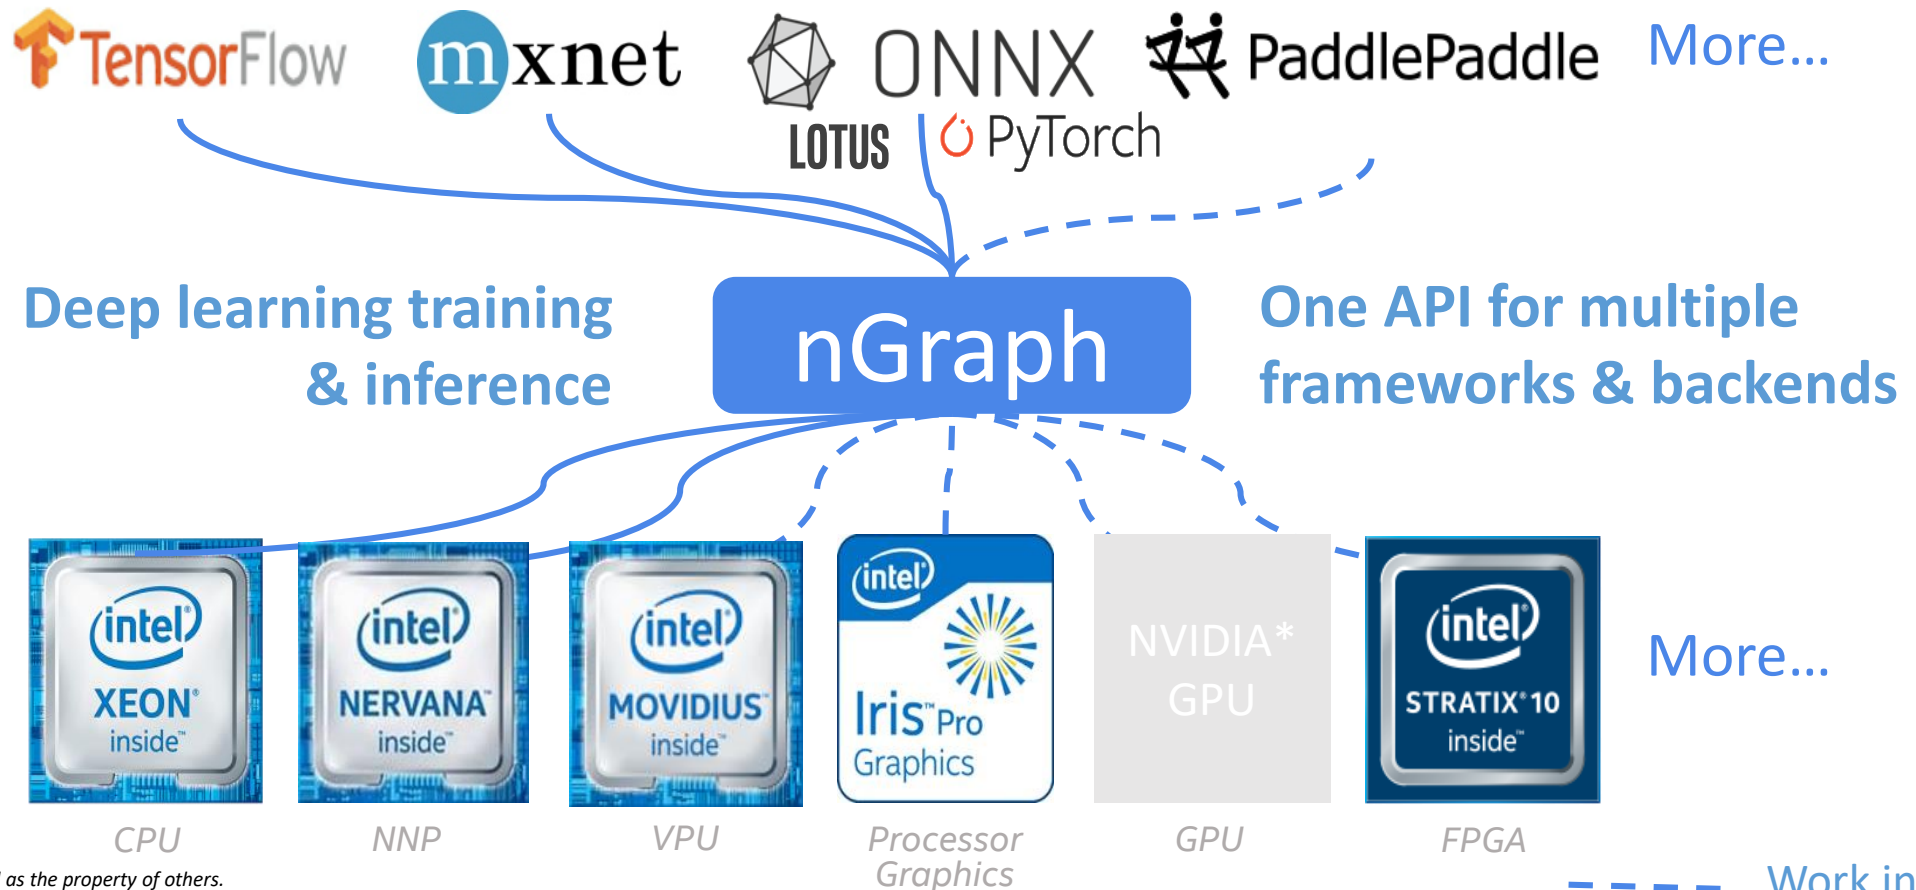
INTEL® INNOVATION DAY

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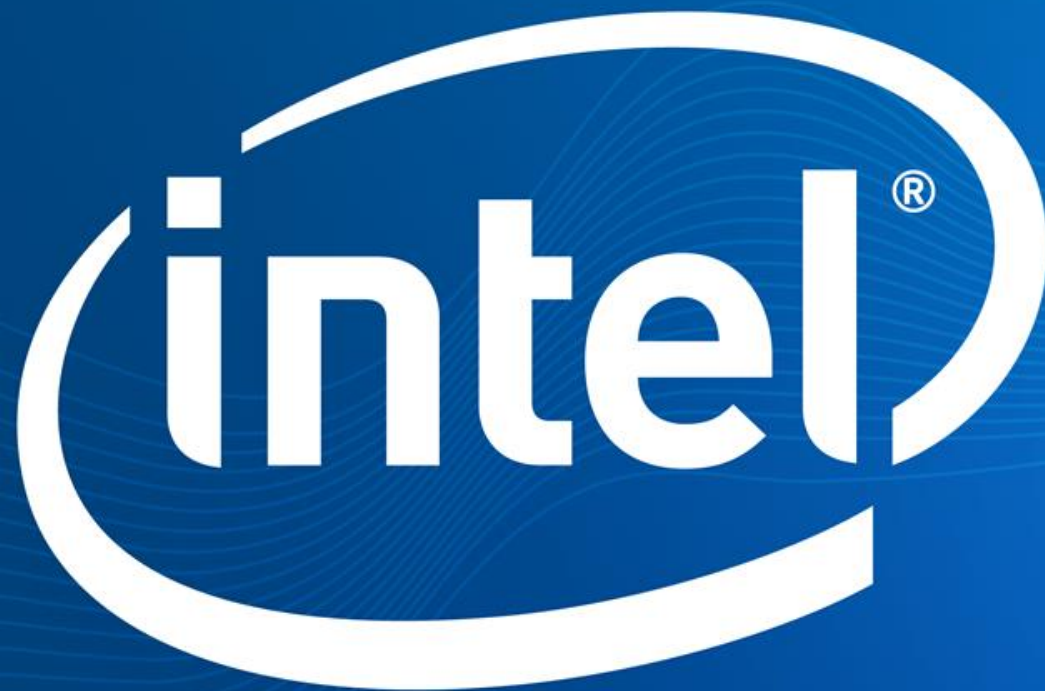OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos

3

Q/A

INTEL® NGRAPH™ COMPILER

Open-source C++ library, compiler & runtime for deep learning enabling flexibility to run models across a variety of frameworks and hardware

TensorFlow   mxnet   ONNX   PaddlePaddle   More...
                       LOTUS PyTorch

Deep learning training & inference

nGraph

One API for multiple frameworks & backends

CPU — XEON inside
NNP — NERVANA inside
VPU — MOVIDIUS inside
Processor Graphics — Iris Pro Graphics
GPU — NVIDIA* GPU
FPGA — STRATIX 10 inside

More...

- - - - Work in progress

*Other names and brands may be claimed as the property of others.
All products, computer systems, dates, and figures are preliminary based on current expectations, and are subject to change without notice.