



INTEL AI PORTFOLIO

Legal Notices and Disclaimers

This presentation is for informational purposes only. INTEL MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at [intel.com](https://www.intel.com).

This sample source code is released under the [Intel Sample Source Code License Agreement](#).

Intel, the Intel logo, neon, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Copyright © 2018, Intel Corporation. All rights reserved.

Outline: What students are expected to learn

- Intel vision for AI on PC
- Introduction to the Intel analytics ecosystem
- Introduction to the Intel hardware for training and inference
- Intel – from hardware to solutions
- Summary



Bring Your AI Vision to Life Using Intel® Comprehensive Portfolio

DATA

Intel analytics ecosystem to get your data ready

SOLUTIONS

Partner ecosystem to facilitate AI in finance, health, retail, industrial & more

TOOLS

Software to accelerate development and deployment of real solutions

HARDWARE

Multi-purpose to purpose-built AI compute from device to cloud

FUTURE

Driving AI forward through R&D, investments and policy

DATA

Intel analytics ecosystem
to get your data ready



011010110110
110101101011
001011010100



INTEGRATE



STORE



PROCESS



MANAGE



ANALYZE

SOURCE(S)?
STRUCTURED?
VOLUME?
DURABILITY?
STREAMING?
LOCALITY?
GOVERNANCE?
OTHER?

Tool for **live streaming data ingestion** from Internet of Things (IOT) sensors in endpoint devices

e.g. Kafka, Sqoop*, MQTT*, WS*, REST*, Flume**

File, block or object-based storage solution given cost, access, volume and performance requirements

e.g. Lustre, IBM* Spectrum* Scale* (GPFS), Dell/EMC* Isilon*, MySQL* (OLTP), Tera*data* (EDW), AWS* S3* (ODS), HDFS* (No-SQL), Hbase* (In-Mem DB)*

Integration, cleaning, normalization and other transformations on batch and/or streaming data

e.g. Hadoop MapReduce*, Apache* Storm*, Beam**

Job scheduling and storage management framework for distributed computation in various domains

e.g. SLURM, PBS*, YARN*, Mesos*, Kubernetes**

Applications in **HPC, Big Data, HPDA, AI** & more that have access to a common compute and data pool

e.g. MPI, SHMEM*, Hadoop*, Spark*, Apache*, Flink*, TensorFlow*, MXNet**

*Other names and brands may be claimed as the property of others. Non-exhaustive list of offerings in each category.

[INTEL.COM/ANALYTICS](https://www.intel.com/analytics)



HARDWARE

Multi-purpose to purpose-built
AI compute from cloud to device



MAINSTREAM

INTENSIVE

DEEP
LEARNING
→ **TRAINING**
→ **INFERENCE**



MOST
OTHER **AI**



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

ONE SIZE DOES NOT FIT ALL



HARDWARE

Multi-purpose to purpose-built
AI compute from device to cloud



END POINT



User-touch end point devices with lower power requirements such as laptops, tablets, smart home devices, drones

EDGE



Small scale data centers, small business IT infrastructure, to few on-premise server racks and workstations

DATA CENTER



Large scale data centers such as public cloud or comms service providers, gov't and academia, large enterprise IT

← Varies to <1ms <5ms <10-40ms ~100ms →

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

AI IS EXPANDING



HARDWARE

Multi-purpose to purpose-built
AI compute from device to cloud



END POINT



IOT SENSORS
(Security, home, retail, industrial...)



Vision & Inference

Speech



SELF-DRIVING VEHICLES



Autonomous Driving



DESKTOP & MOBILITY



SOC

M.2 Card

Display, Video, AR/VR, Gestures, Speech

EDGE

SERVERS, APPLIANCES & GATEWAYS



Most Use Cases

+ Special Purpose



Dedicated Media & Vision Inference



Latency-Bound Inference



Basic Inference, Media & Vision

DATA CENTER

SERVERS & APPLIANCES



Most Use Cases



NNP-L / NNP-I

Most Intensive Use Cases

+ Special Purpose



Flexible & Memory Bandwidth-Bound Use Cases

Varies to <1ms <5ms <10-40ms ~100ms

¹GNA=Gaussian Neural Accelerator
All products, computer systems, dates, and figures are preliminary based on current expectations, and are subject to change without notice. Images are examples of intended applications but not an exhaustive list.

ONE SIZE DOES NOT FIT ALL

Intel Confidential





SOLUTIONS

Solution Architects

TOOLKITS

App Developers

LIBRARIES

Data Scientists

FOUNDATION

Library Developers

HARDWARE

IT System Architects

AI
RN
TT
IE
FL
IL
CI
G
AE
LN
CE

ARTIFICIAL INTELLIGENCE

AI Solutions Catalog
(Public & Internal)



Platforms



Finance



Healthcare



Energy



Industrial



Transport



Retail



Home



More...

DEEP LEARNING DEPLOYMENT

OpenVINO™ †

Open Visual Inference & Neural Network Optimization toolkit for inference deployment on CPU, processor graphics, FPGA & VPU using TF, Caffe* & MXNet*

Intel® Movidius™ SDK

Optimized inference deployment for all Intel® Movidius™ VPUs using TensorFlow* & Caffe*

DEEP LEARNING COMING SOON! Intel® Deep Learning Studio†

Open-source tool to compress deep learning development cycle

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*](#)



[PaddlePaddle*](#)

ANALYTICS, MACHINE & DEEP LEARNING PRIMITIVES

Python

Intel distribution optimized for machine learning

DAAL

Intel® Data Analytics Acceleration Library (for machine learning)

MKL-DNN

Open-source deep neural network functions for CPU, processor graphics

cLDNN

DEEP LEARNING GRAPH COMPILER

Intel® nGraph™ Compiler (Alpha)

Open-sourced compiler for deep learning model computations optimized for multiple devices (CPU, GPU, NNP) using multiple frameworks (TF, MXNet, ONNX)

AI FOUNDATION



Data Center
Edge
Device



NP L-1000



Inference

DEEP LEARNING ACCELERATORS

† Formerly the Intel® Computer Vision SDK
*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.



SUMMARY

Summary

- Intel has technology in Data Center, Cloud, and PC to support AI.

1. Hardware

Capabilities highlighted:

- HW CPU – i7 Core 8th Gen, i7-8550U (CFL)
- HW CPU – i7 Core 8th Gen i7-8565U (WKL)
- Intel Integrated GPU – HD graphics
- Camera, Movidius Neural Compute Stick
- WinML via Microsoft Windows 10 – 64 bit
- OpenVINO™ Toolkit, Intel Dist. of Python
- IDE: Visual Studio Community Ed. 2017
- TensorFlow with Tutorials, Labs

2. Tools

3. Solutions

- Intel is positioning the PC as a critical AI engine at the edge.

