Machine Learning & Artificial intelligence and Cyber Security

Jeff Towle
Intel Corporation
Sr. Cloud Security Architect
Next Wave Cloud Computing Group

Legal notices & disclaimers

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.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer. No computer system can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit http://www.intel.com/performance.

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.

Statements in this document that refer to Intel's plans and expectations for the quarter, the year, and the future, are forward-looking statements that involve a number of risks and uncertainties. A detailed discussion of the factors that could affect Intel's results and plans is included in Intel's SEC filings, including the annual report on Form 10-K.

The products described 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.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

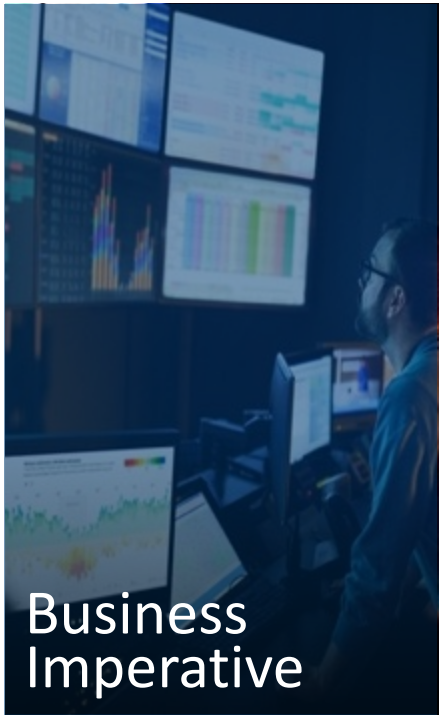
Business
Imperative


What is AI?


AI & ML Cyber
Security

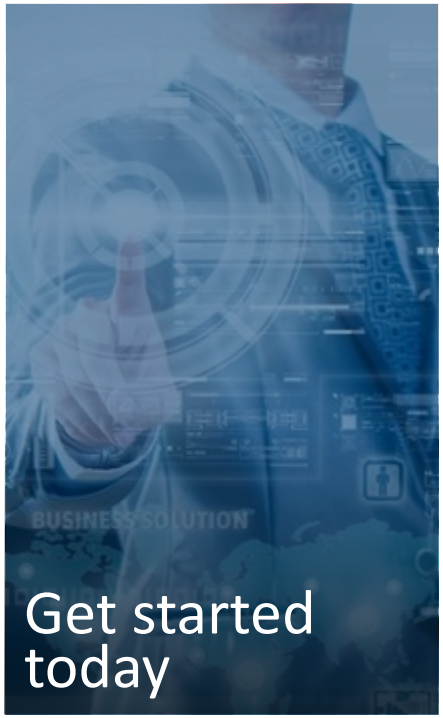
Get started
today

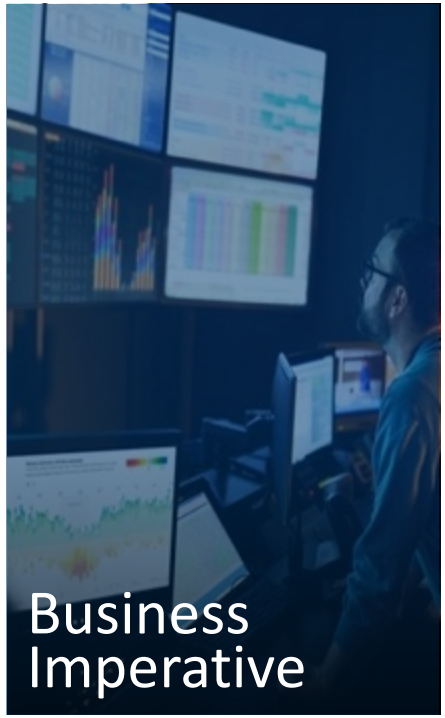
Business
Imperative


What is AI?

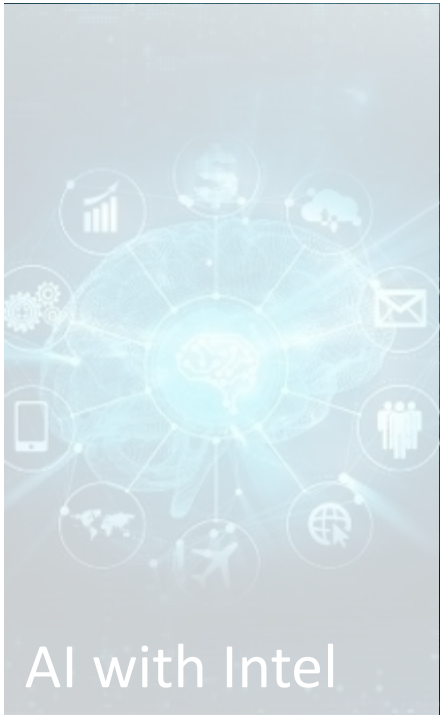
AI with Intel

Get started
today

AI will transform

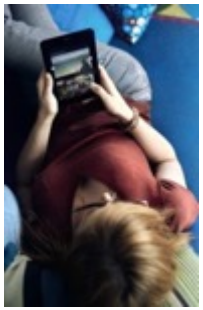



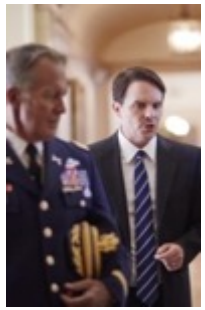



Consumer	Health	Finance	Retail	Government	Energy	Transport	Industrial	Other
Smart Assistants	Enhanced Diagnostics	Algorithmic Trading	Support	Defense	Oil & Gas Exploration	In-Vehicle Experience	Factory Automation	Advertising
Chatbots	Drug Discovery	Fraud Detection	Experience	Data Insights	Smart Grid	Automated Driving	Predictive Maintenance	Education
Search	Patient Care	Research	Marketing	Safety & Security	Operational Improvement	Aerospace	Precision Agriculture	Gaming
Personalization	Research	Personal Finance	Merchandising	Resident Engagement	Conservation	Shipping	Field Automation	Professional & IT Services
Augmented Reality	Sensory Aids	Risk Mitigation	Loyalty	Smarter Cities		Search & Rescue		Telco/Media
Robots			Supply Chain					Sports
			Security					

Source: Intel forecast

Ai adoption is nascent

Source: Gartner Says Nearly Half of CIOs Are Planning to Deploy Artificial Intelligence. February 2018 (https://www.gartner.com/newsroom/id/3856163)

The path to deeper insight

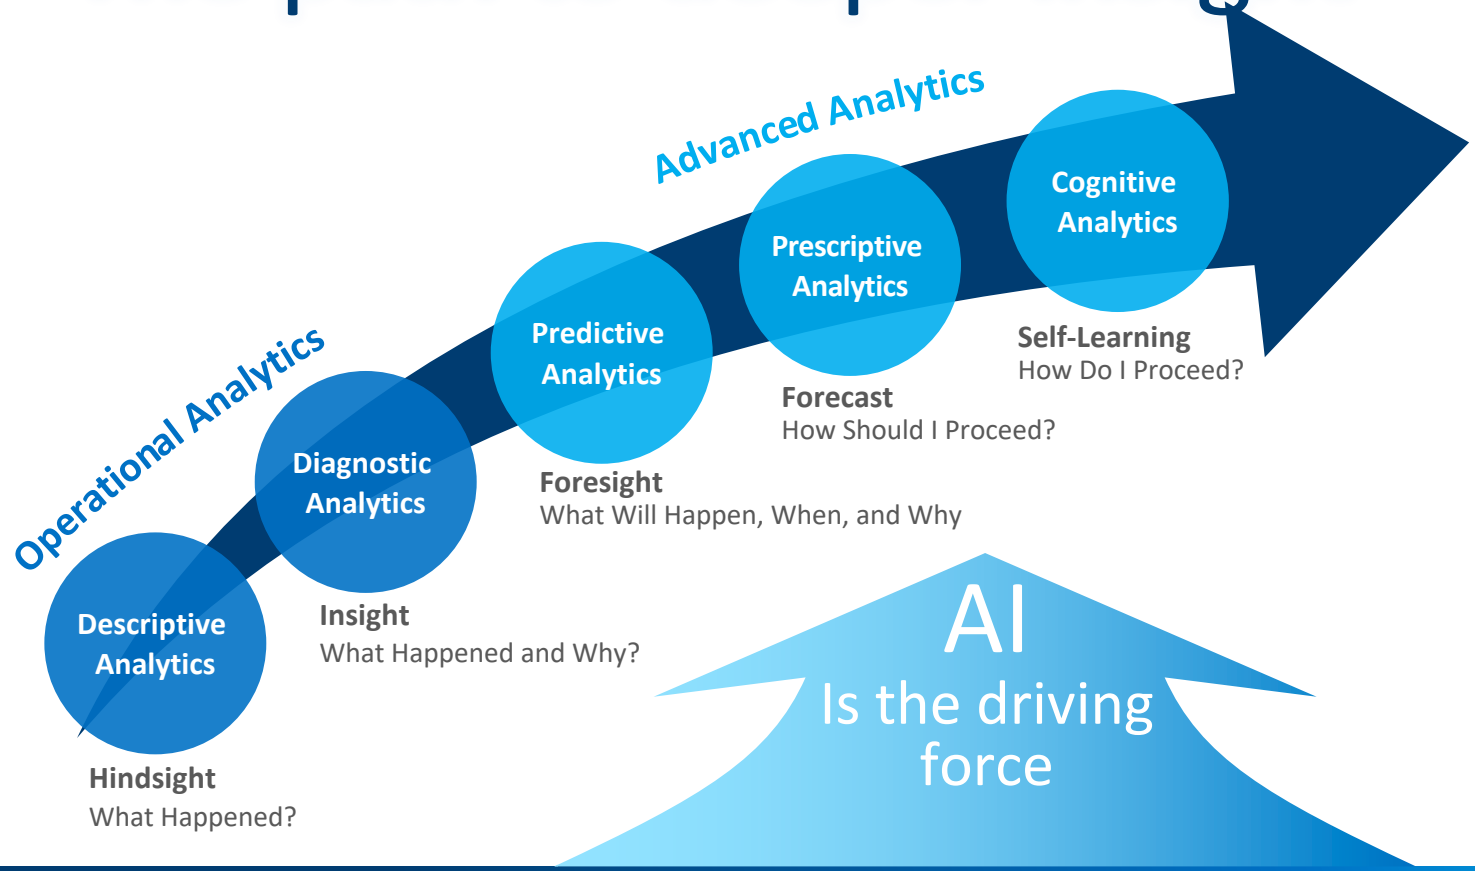

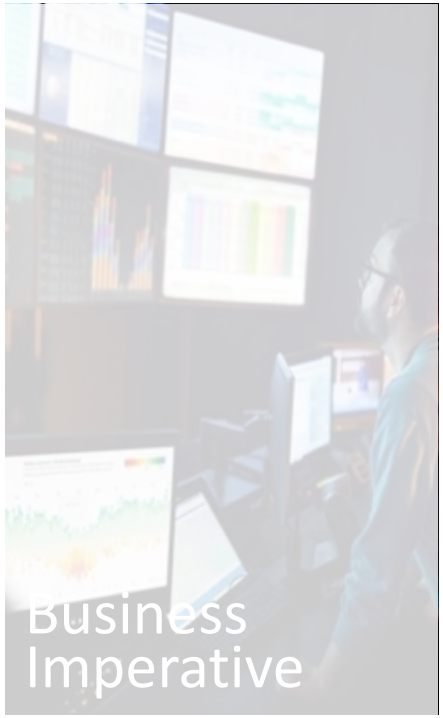
Business
Imperative


What is AI?

AI with Intel

Get started
today

Artificial Intelligence

is the ability of machines to learn from experience, without explicit programming, in order to perform cognitive functions associated with the human mind

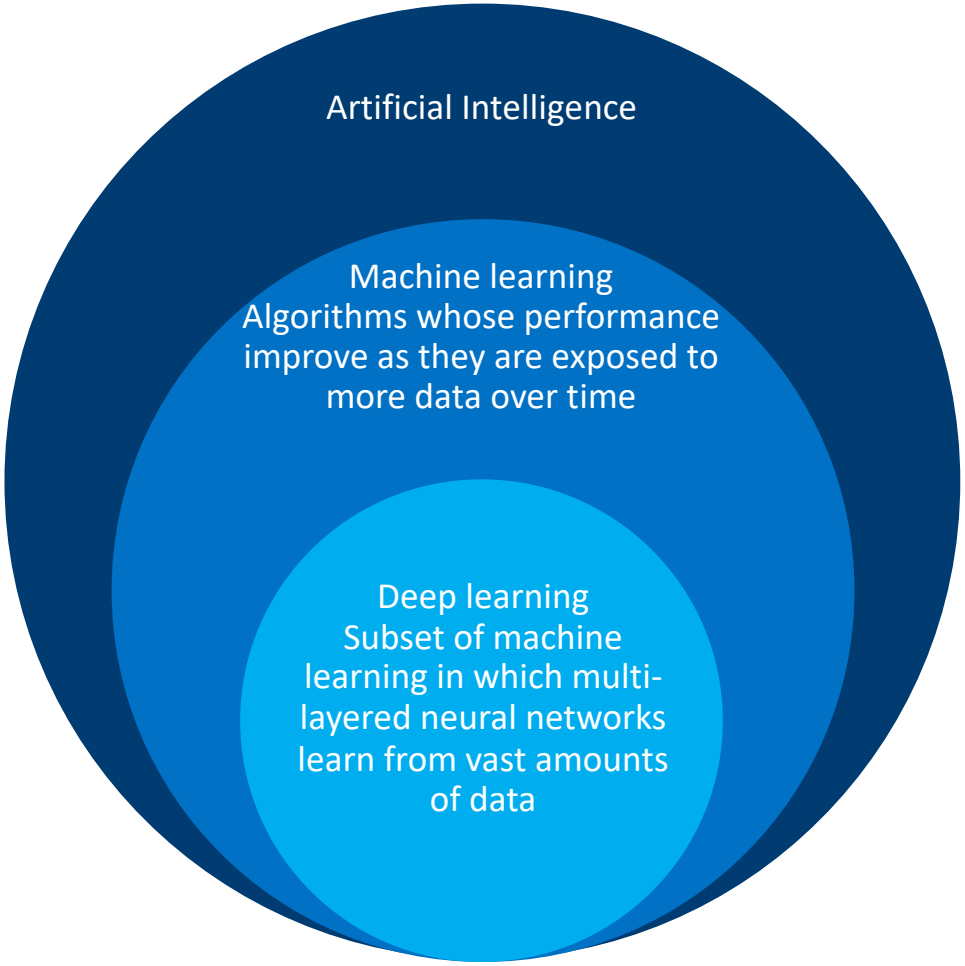

The AI lifecycle

1. Define the Challenge

2. Approach
Team breaks down the defined business problem into workable steps to translate the right data to achieve results

3. Expertise
A team of management sponsors, data scientists, data engineers, solution architects, and domain experts identifies the right data and works to translate the data to achieve results

4. Philosophy
Team embraces fail-fast continuous improvement practices to evaluate their success in translating data to achieve results

5. Source Data
Team understands and obtains the right data that explains the business problem to achieve results

6. Infrastructure
Organization secures hardware and software infrastructure that supports data processing in a timely manner

7. Organization
Organization embraces data insights, sponsors properly resourced teams, and prioritizes analytic development work

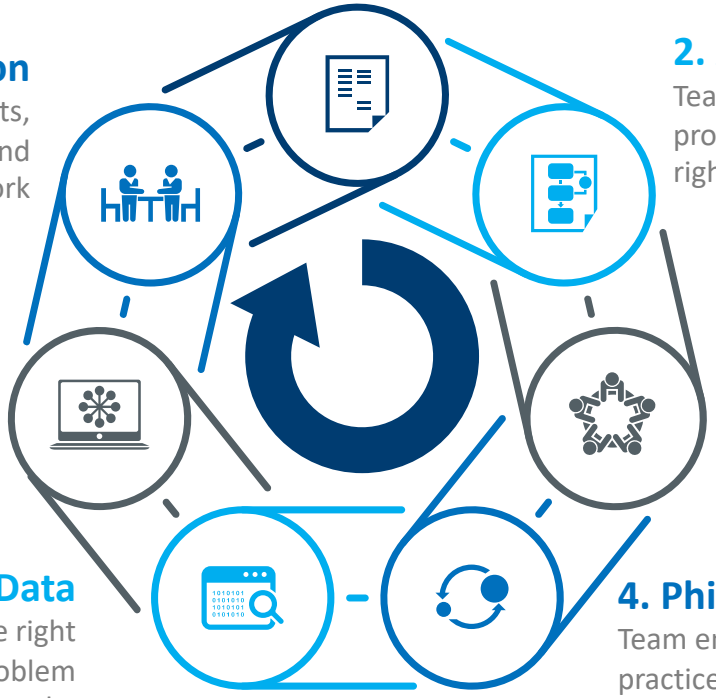

Machine vs. Deep learning

Machine Learning

How do you engineer the best features?

$N \times N$

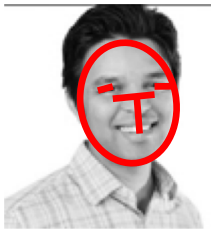

$(f_1, f_2, \ldots, f_K)$

Roundness of face
Dist between eyes
Nose width
Eye socket depth
Cheek bone structure
Jaw line length
…etc.

CLASSIFIER ALGORITHM

SVM
Random Forest
Naïve Bayes
Decision Trees
Logistic Regression
Ensemble methods

Arjun

Deep Learning

How do you guide the model to find the best features?

$N \times N$

NEURAL NETWORK

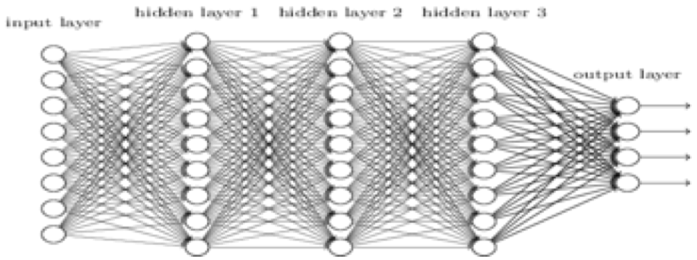

Arjun

Deep learning Basics

Training

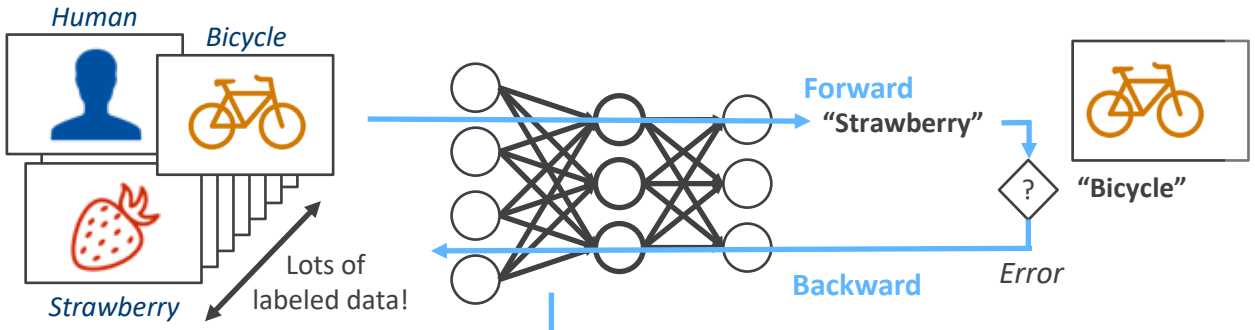

Inference

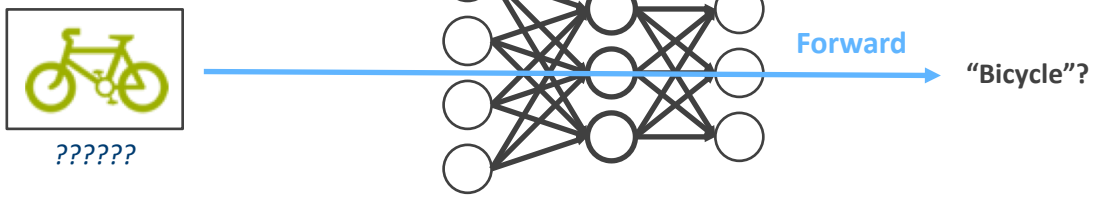

Did you know?

Training with a large data set AND deep (many layered) neural network often leads to the highest accuracy inference

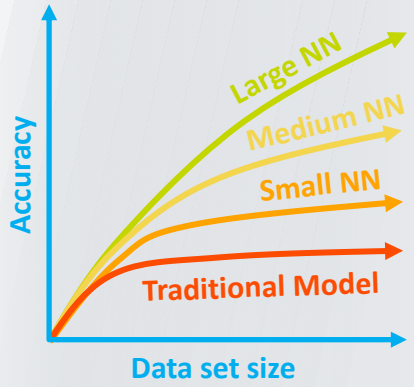

Deep learning breakthroughs

Machines able to meet or exceed human image & speech recognition

Image recognition

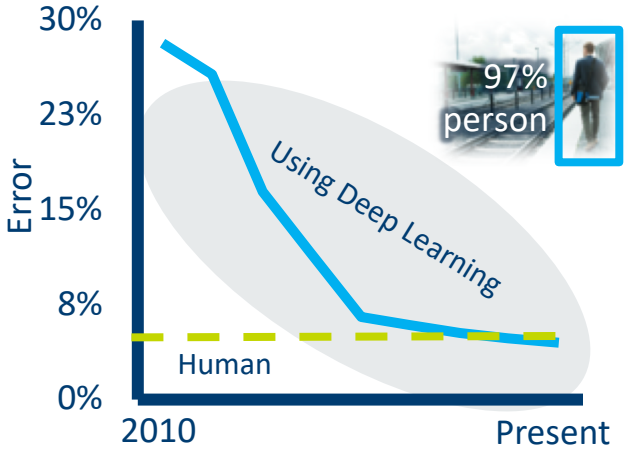

Speech recognition

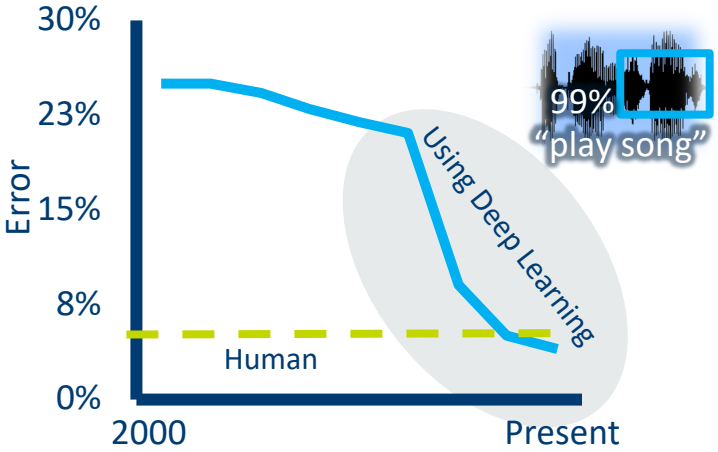

e.g.

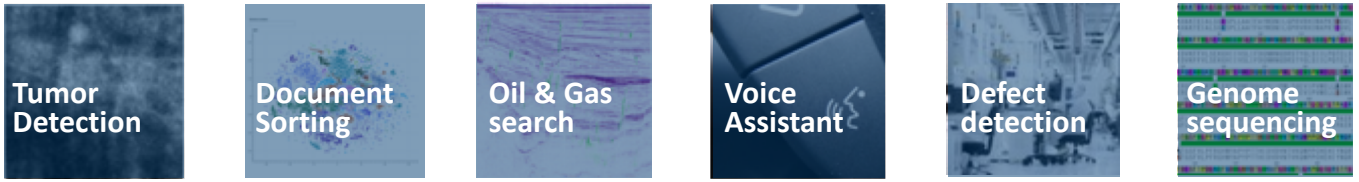

Source: ILSVRC ImageNet winning entry classification error rate each year 2010-2016 (Left), https://www.microsoft.com/en-us/research/blog/microsoft-researchers-achieve-new-conversational-speech-recognition-milestone/ (Right)

AI customer example

Intel works with customers across the entire AI lifecycle

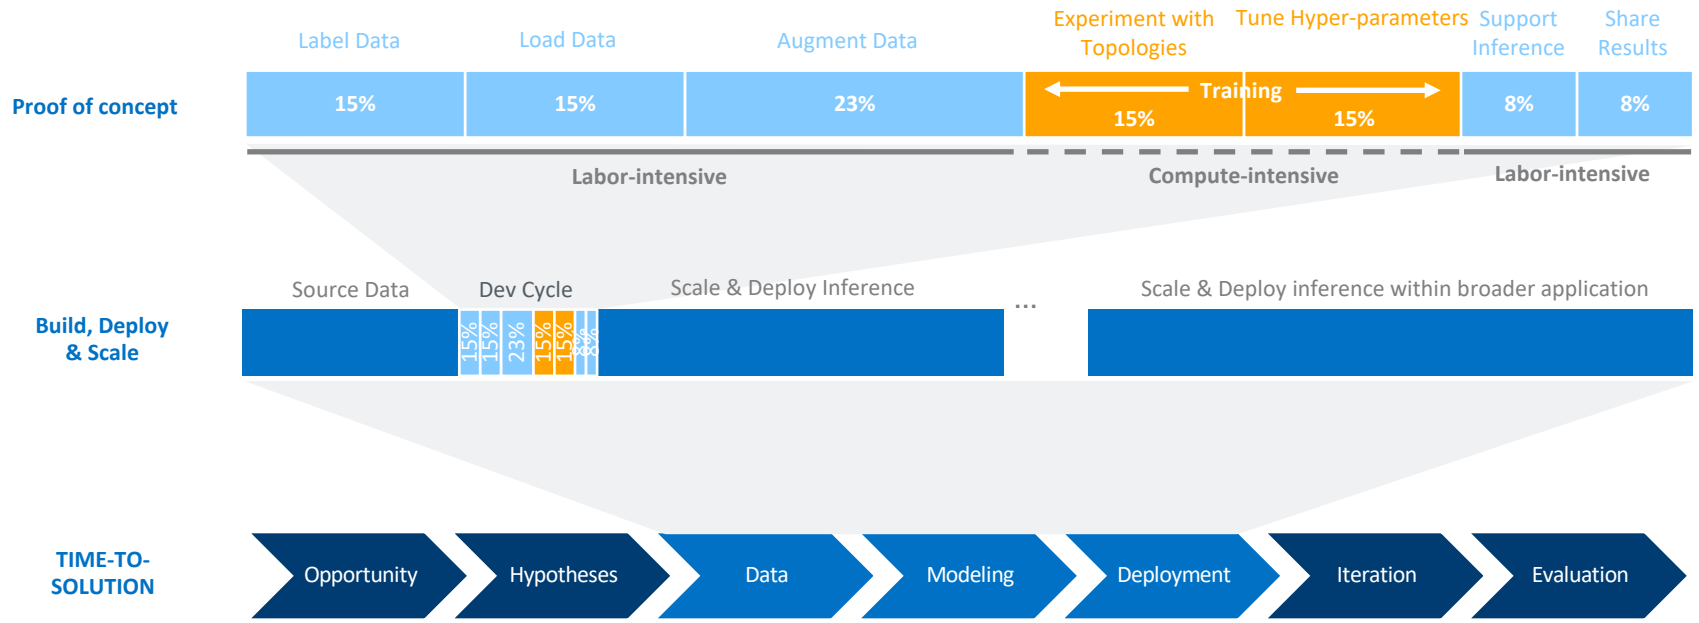

AI customer example

Intel helps customers deploy & scale real AI solutions

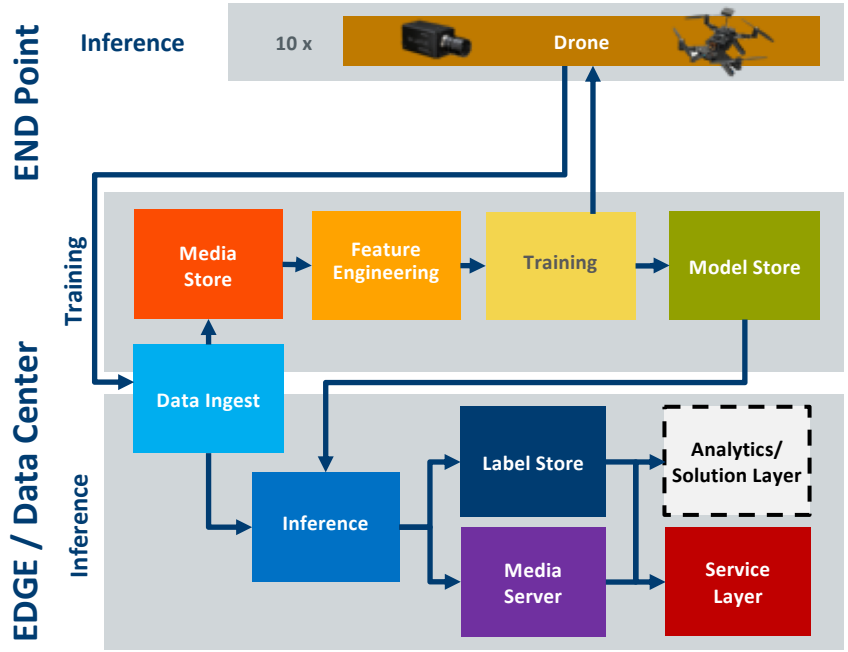

Software: *OpenVino™* *Intel® Movidius™ SDK* *TensorFlow** *Intel® MKL-DNN*

Drone:
- Inference
- 10 Drones
- Inference

Real-time object detection

Per Drone
- 1x Intel® Core™ processor
- 1x Intel® Movidius™ VPU

Media Server:
- Media Store
- Media Store
- Media Store

110 Nodes
- *8 TB/day per camera*
- *10 cameras*
- *3x replication*
- *1-year video retention*
- *4 mgmt nodes*

- Media Store
- Media Store
- Media Store

Per Node
- 1x 2S 61xx
- 20x 4TB SSD

Multi-Purpose Cluster:
- Data Ingestion
- Data Ingestion
- Data Ingestion
- Data Ingestion
- Inference
- Inference
- Inference
- Inference
- Tagged Datasets
- Tagged Datasets
- Service Layer
- Service Layer
- Service Layer
- Media Server
- Media Server
- Media Server

4 nodes *One ingestion per day, one-day retention*

4 nodes *20M frames per day*

2 nodes *Infrequent op*

3 nodes *Simultaneous users*

3 nodes *10k clips stored*

Per Node
- 1x 2S 61xx
- 20x 4TB SSD

Data Storage:
- Model Store
- Model Store
- Model Store
- Model Store
- Label Store
- Label Store
- Label Store
- Label Store

4 nodes *1-year of history*

4 nodes *Labels for 20M frames /day*

Per Node
- 1x 2S 81xx
- 5x 4TB SSD

Training:
- Training

16 nodes *<10 hours TTT*
- Training

Per Node
- 1x 2S 81xx
- 1x 4TB SSD

Source: Intel customer engagement
*Other names and brands may be claimed as the property of others

Intel® OpenVINO: Funny Name, Great Strategy*

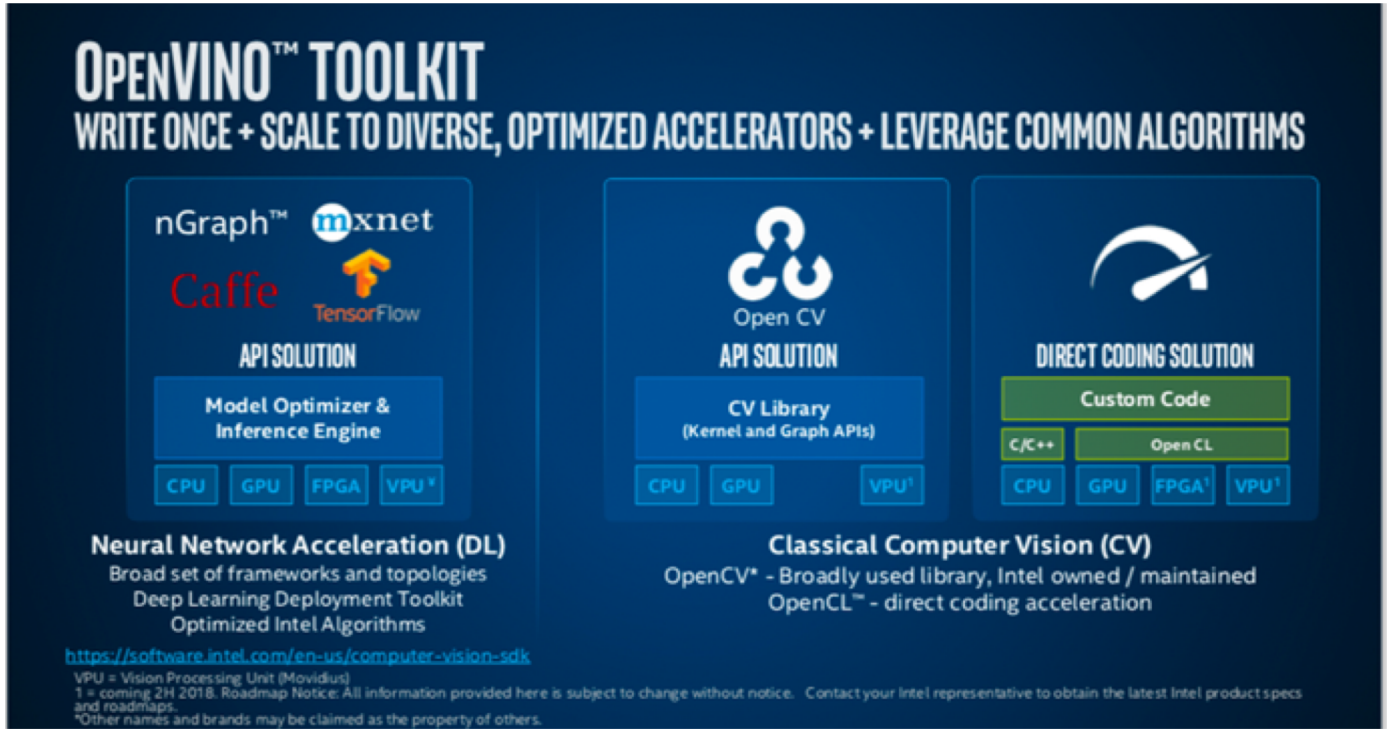

https://www.forbes.com/sites/moorinsights/2018/05/22/intel-openvino-funny-name-great-strategy/#5c3d9f846301

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

CELERON | ATOM | CORE | XEON | ARRIA 10 | MOVIDIUS | IRIS Pro

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.

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® Deep Learning Deployment Toolkit (DLDT)

Take Full Advantage of the Power of Intel® Architecture for Deep Learning

Model Optimizer

- **What it is**: Preparation step -> imports trained models
- **Why important**: Optimizes for performance/space with conservative topology transformations; biggest boost is from conversion to data types matching hardware.

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.

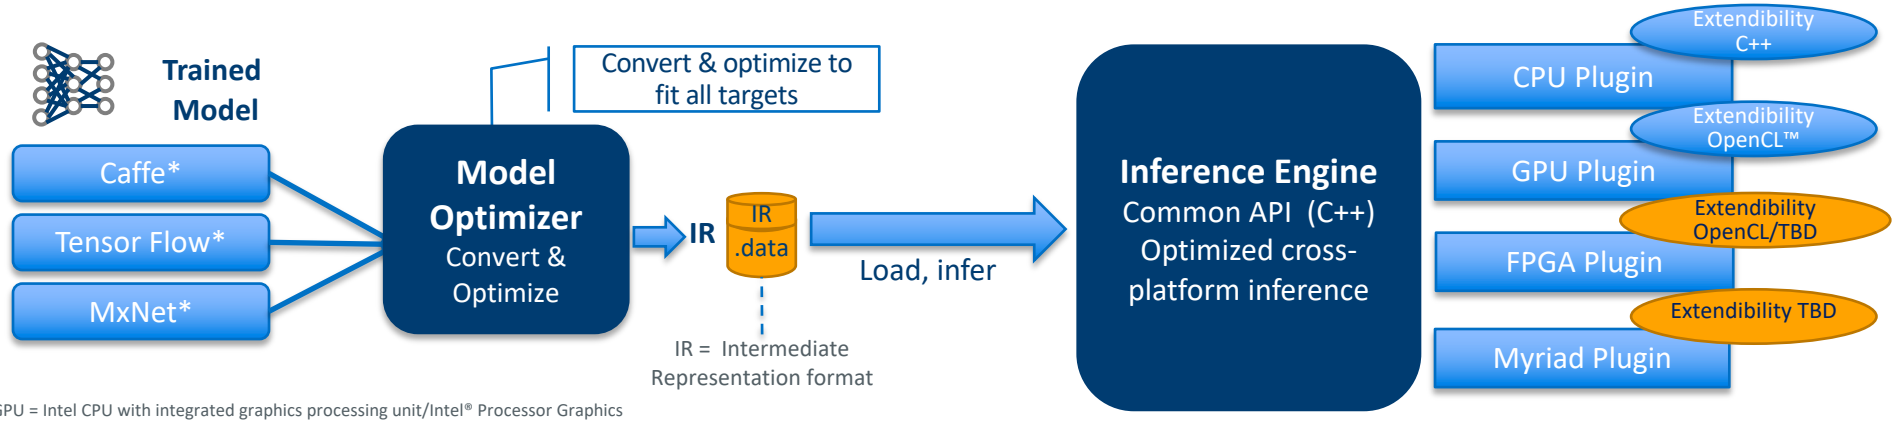

GPU = Intel CPU with integrated graphics processing unit/Intel® Processor Graphics

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

Intel® Xeon® scalable processors

Continued Deep Learning Optimization to Deliver Increased Performance

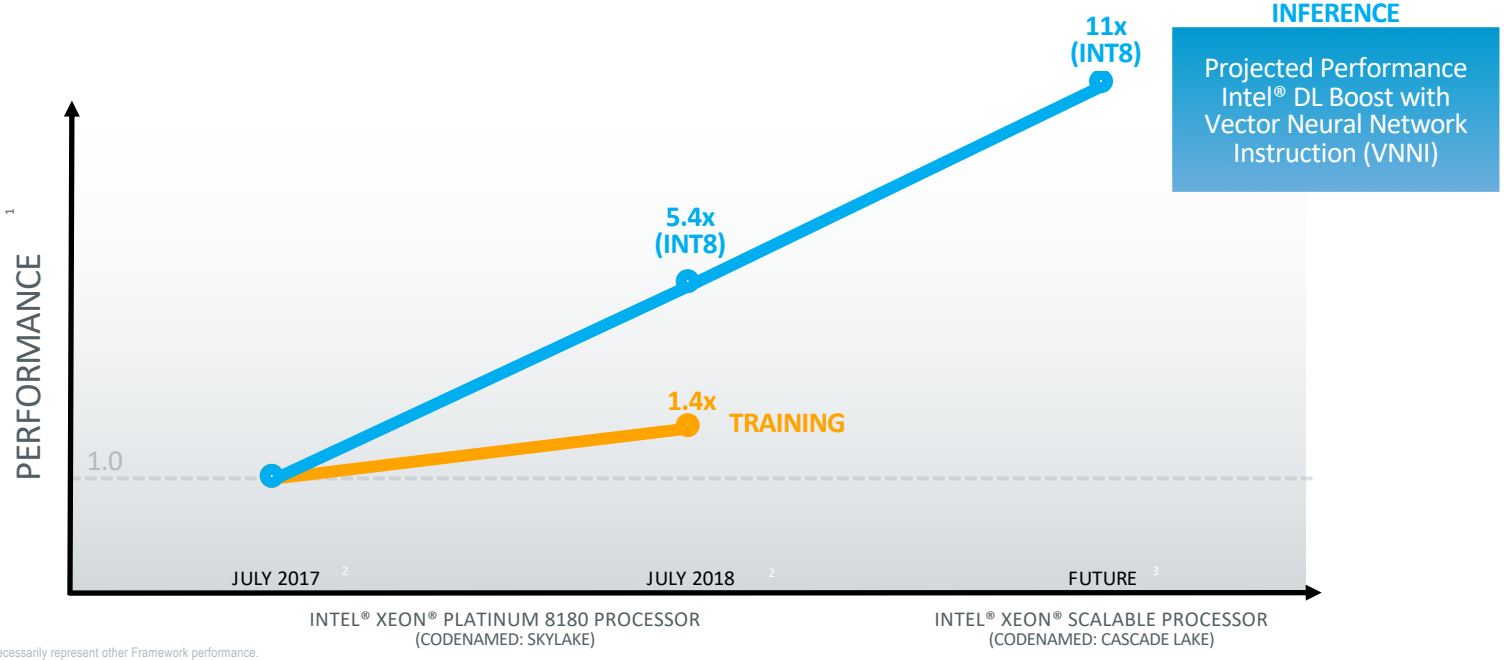

[1] Intel® Optimization for Caffe Resnet-50 performance does not necessarily represent other Framework performance.

[2] Based on Intel internal testing: 1X (7/11/2017), 2.8X (1/19/2018), 1.4x (8/2/2018) and 5.4X (7/26/2018) performance improvement based on Intel® Optimization for Café Resnet-50 inference throughput performance on Intel® Xeon® Scalable Processor. See Configuration Details Slide #19

Performance results are based on testing as of 7/11/2017(1x), 1/19/2018(2.8x), 8/2/2018 (1.4x) & 7/26/2018(5.4) and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

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. For more complete information visit: http://www.intel.com/performance.

Intel® Xeon® scalable processors

Continued Deep Learning Optimization to Deliver Increased Performance

These VNNI instructions are all done within the AVX-512 vector units. Here is how the 8-bit VNNI instruction works compared to Skylake:

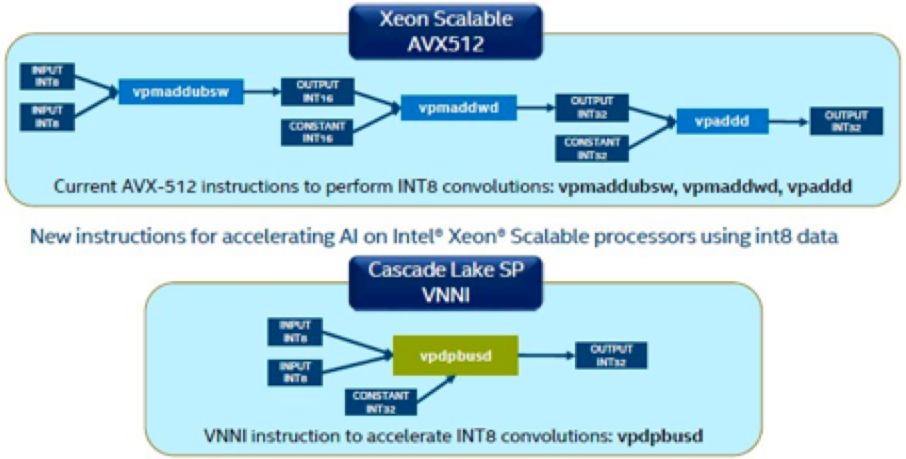

And here is how the 16-bit VNNI instruction works:

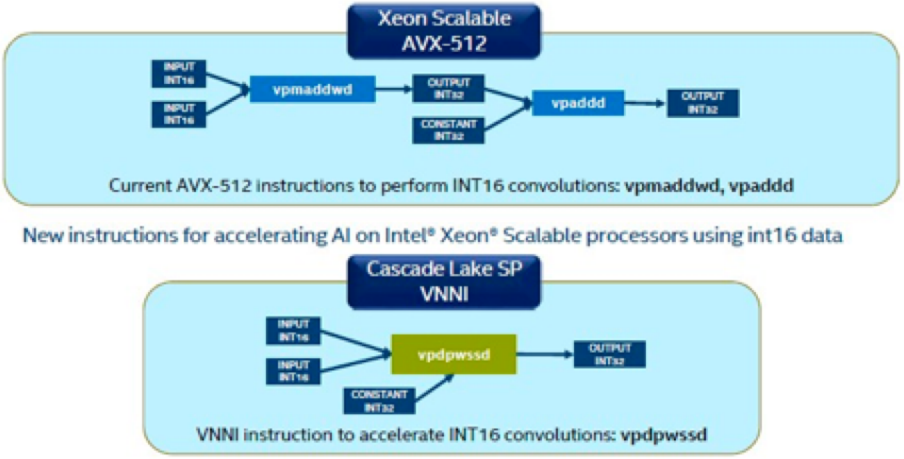

Intel® Xeon® scalable processors

Continued Deep Learning Optimization to Deliver Increased Performance

VNNI Per Core Throughput

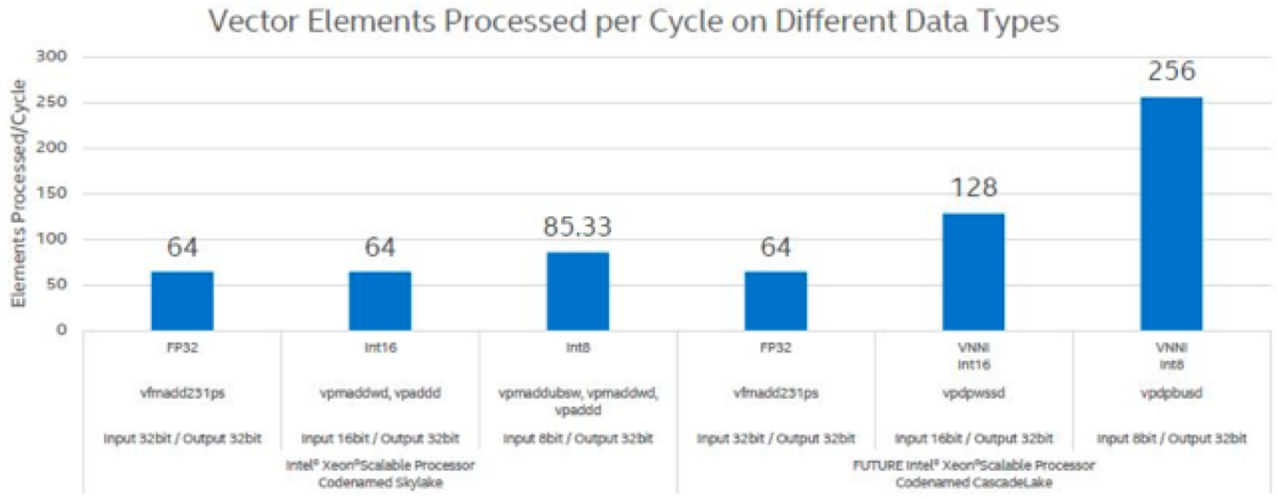

Business
Imperative

What is AI?


AI & ML Cyber
Security

Get started
today

The deluge of data

Daily By 2020

Average internet user	1.5 GB
Autonomous vehicle	4 TB
CONNECTED AIRPLANE	5 TB
Smart Factory	1 PB
Cloud video Provider	750 pB

Source: Amalgamation of analyst data and Intel analysis.

Cisco Tetration

Security, Compliance and Root Cause Analytics

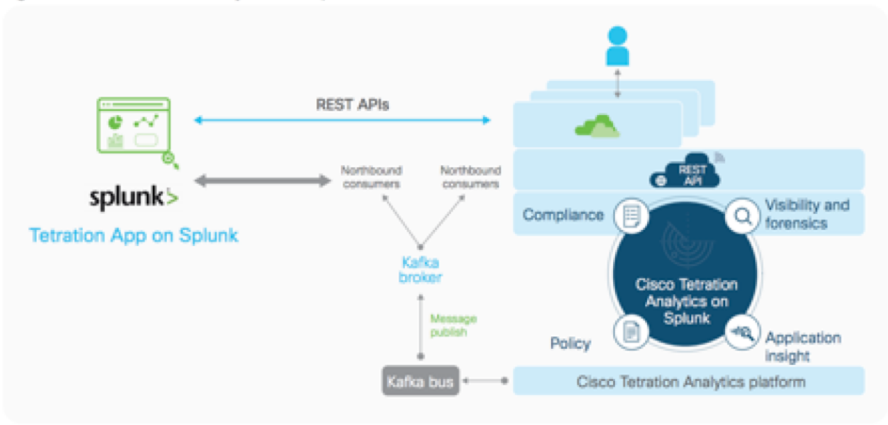

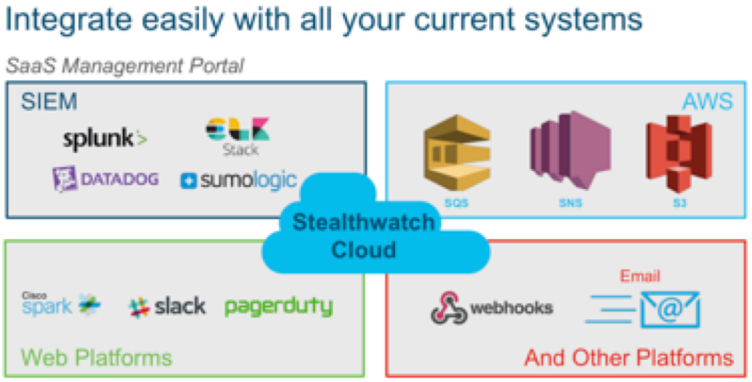

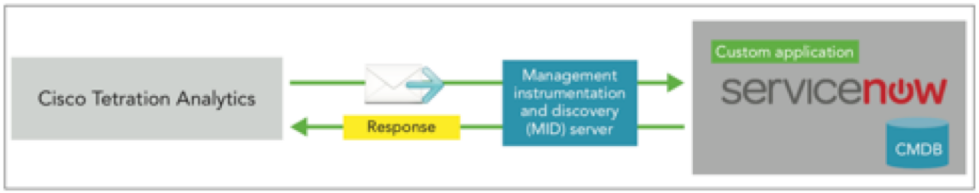

Figure 1. ServiceNow and Cisco Tetration Analytics Solution

https://www.cisco.com/c/dam/m/en_id/dc-innovation/2018/securing-your-dc-arief.pdf

DARKTRACE

The Enterprise Immune System

Detects and fights cyber-threats in real time

- No rules or signatures
- Machine learning and AI
- Real-time threat detection and autonomous response

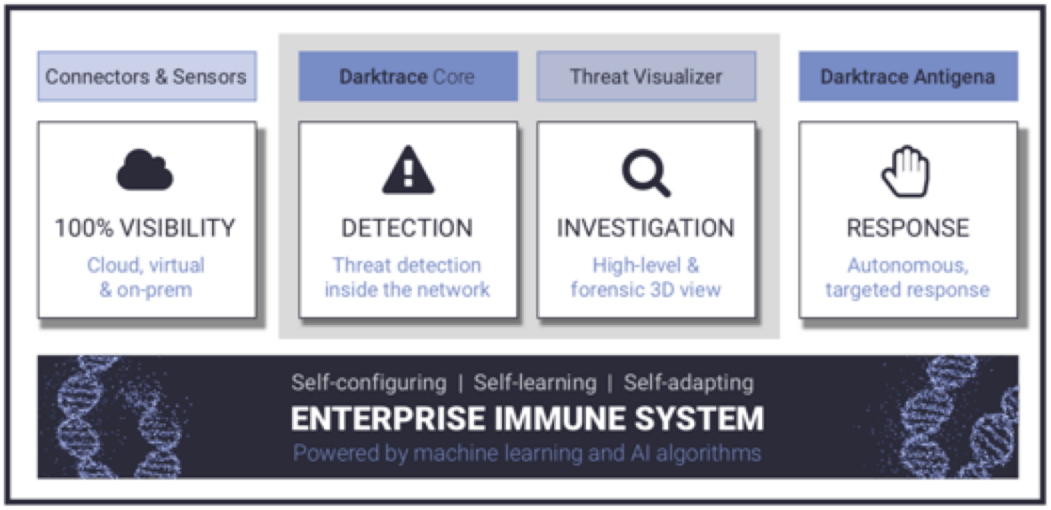

The Enterprise Immune System Version 3 is a self-learning technology platform that protects your entire digital infrastructure, providing 100% visibility into every aspect of your organization including physical networks, cloud, virtualized environments, IoT, and industrial control systems.

amazon web services

sqrrl

Threat Hunting Platform discovers APT's faster, reduce root cause analysis times by:

- Link analysis
- Machine learning algorithms
- Multi-petabyte scalability capabilities into an
- Integrated solution.

harvest.ai

- AI algorithms to protect data stored in cloud services. Prevent cyber attacks and exposure of important documents and information
- "Automatically" identify risk to the business of data that is being exposed or shared outside the organization and remediate based on policies in near real-time,"

Guard Duty

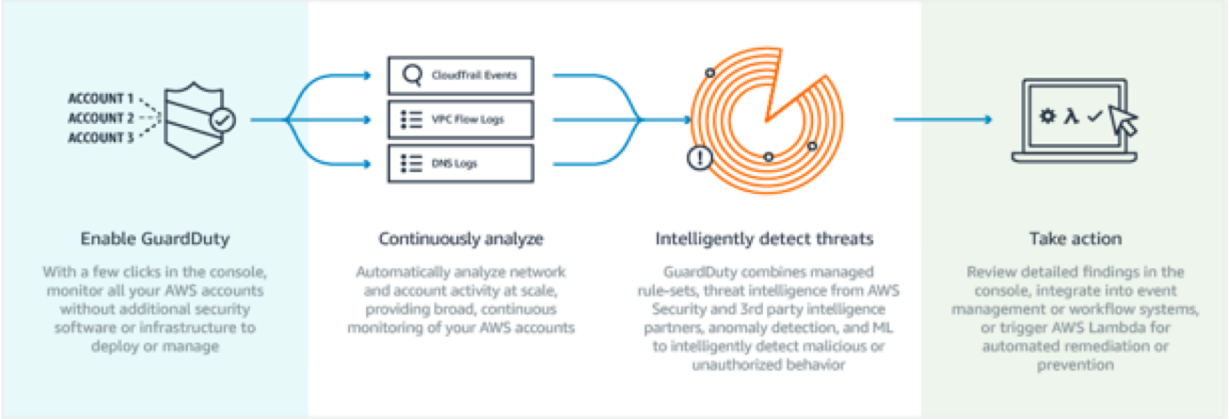

SAS® Viya®

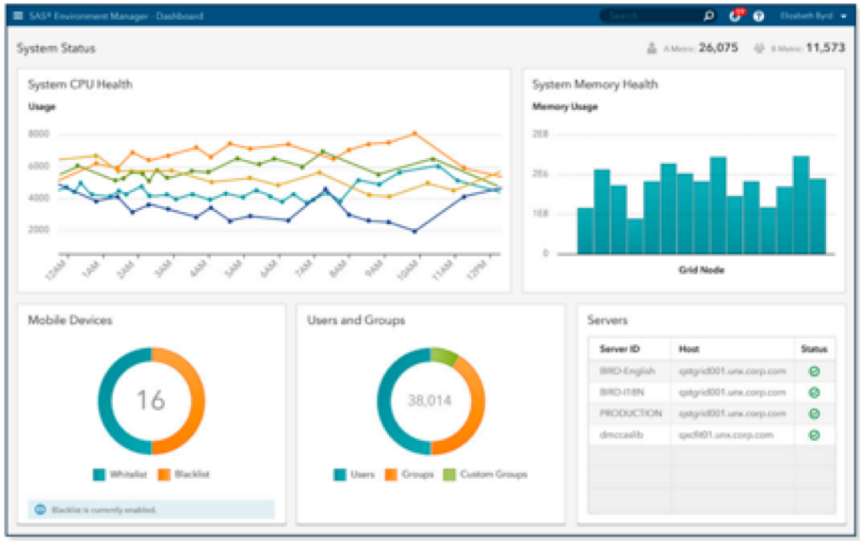
System Status
26,075
11,573
System CPU Health
Usage
System Memory Health
Memory Usage
Grid Node
Mobile Devices
16
Whitelist
Blacklist
Users and Groups
38,014
Users
Groups
Custom Groups
Servers
Server ID
Host
Status

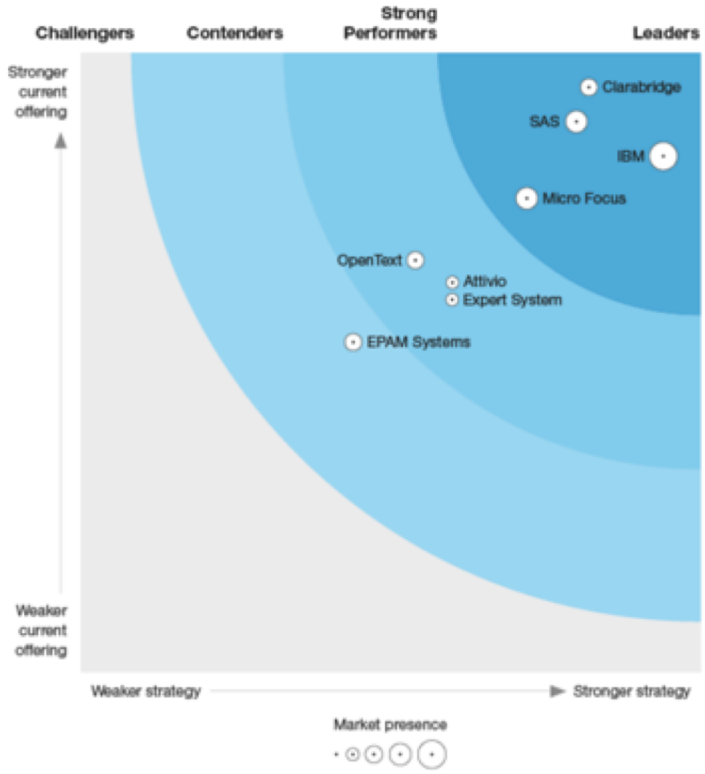
THE FORRESTER WAVE™
AI-Based Text Analytics Platforms
Q2 2018
Challengers
Contenders
Strong Performers
Leaders
Stronger current offering
Weaker current offering
Clarabridge
SAS
IBM
Micro Focus
OpenText
Attivio
Expert System
EPAM Systems
Weaker strategy
Stronger strategy
Market presence

Cylance
Jask I/O
Cybraics
THreatQuotient
Apache Spot
IronNet Security

Apache Spot (Incubating)

A Community Approach to Fighting Cyber Threats

Apache Spot (Incubating)

A Community Approach to Fighting Cyber Threats

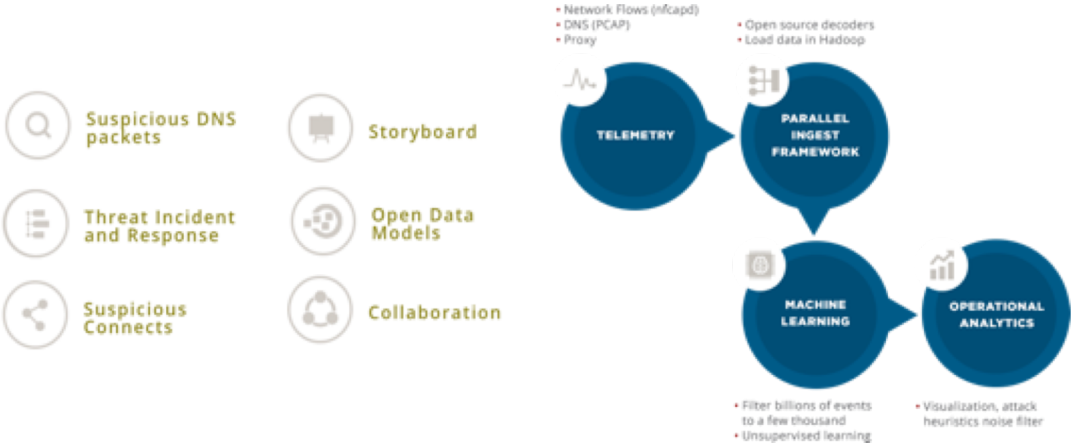

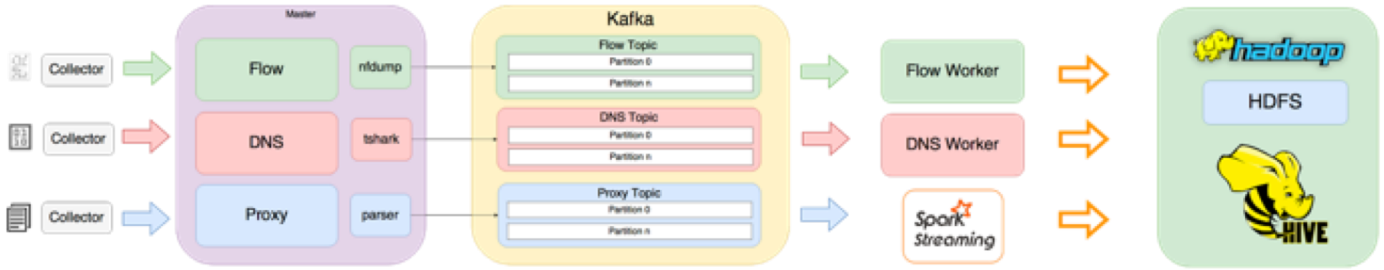

http://spot.incubator.apache.org/

HORTONWORKS®

APACHE METRON

https://hortonworks.com/apache/metron/

THREATQUOTIENT™

THREATQ SOLUTION ARCHITECTURE

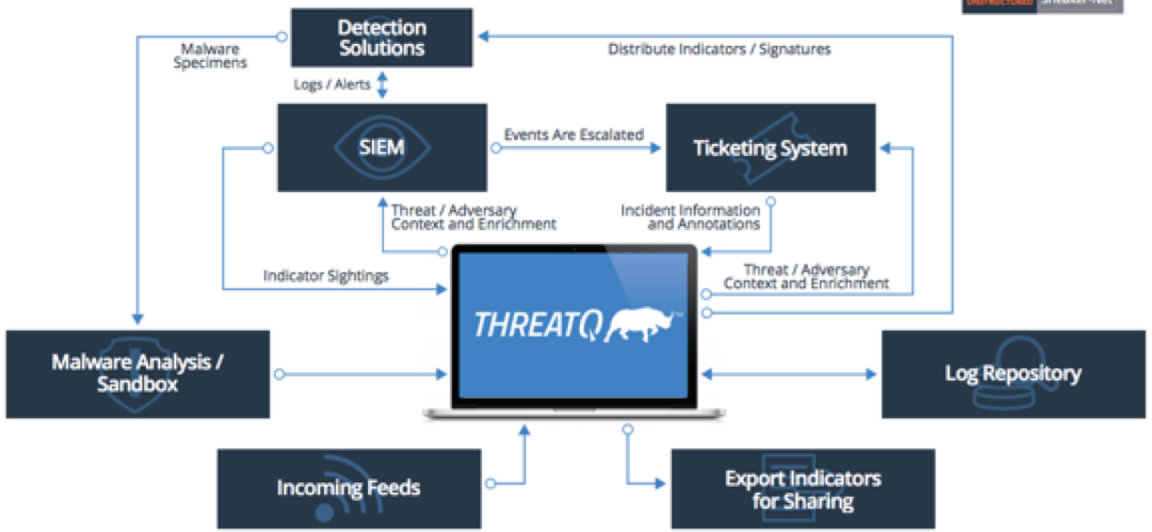

Business
Imperative

What is AI?

AI with Intel

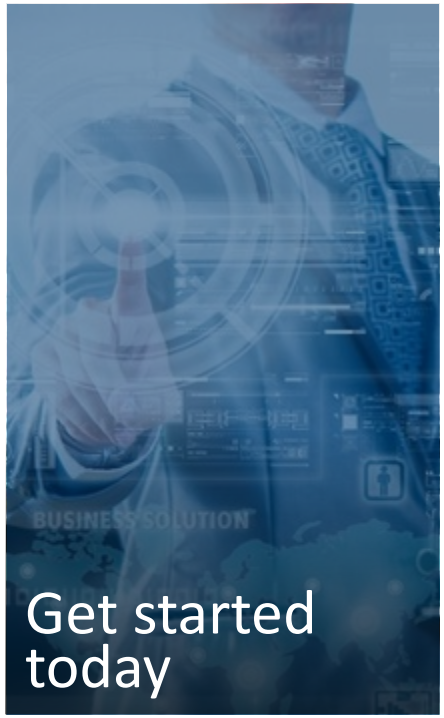
Get started
today

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

Most Intensive Use Cases

Special Purpose

Flexible & Memory Bandwidth-Bound Use Cases

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

[1]GNA=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

ARTIFICIAL INTELLIGENCE

Solutions
Solution Architects

TOOLKITS
App Developers

libraries
Data Scientists

foundation
Library Developers

Hardware
IT System Architects

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*

Coming soon

ANALYTICS, MACHINE & DEEP LEARNING PRIMITIVES

Python
Intel distribution optimized for machine learning

DAAL
Intel® Data Analytics Acceleration Library (for machine learning)

MKL-DNN **clDNN**
Open-source deep neural network functions for CPU, processor graphics

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

DEEP LEARNING ACCELERATORS

Coming 2019

NNP L-1000

Inference

† 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.

HARDWARE

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

intel AI

Mainstrea **intensive**

Training

Deep Learning

+

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

Future

Driving AI forward through R&D, investments and policy

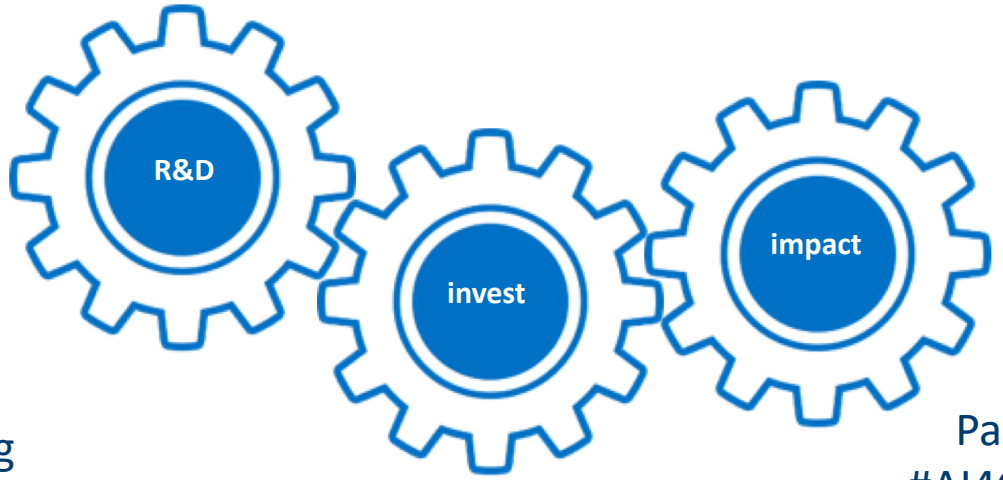

- ✓ Image/Video/Audio
- ✓ Natural Language
- ✓ Autonomous Driving
- ✓ Reinforcement Learning
- ✓ Adversarial Learning
- ✓ Limited Precision/Sparsity
- ✓ More…

Over $1B invested in AI innovators[1]

Partnering with #AI4GOOD to enrich the lives of every person on Earth through AI

[1]Source: https://newsroom.intel.com/editorials/intel-invests-1-billion-ai-ecosystem-fuel-adoption-product-innovation/

Intel® AI academy

For developers, students, instructors and startups

Get smarter using online tutorials, webinars, student kits and support forums

Educate others using available course materials, hands-on labs, and more

Get 4-weeks FREE access to the Intel® AI DevCloud, use your existing Intel® Xeon® Processor-based cluster, or use a public cloud service

Showcase your innovation at industry & academic events and online via the Intel AI community forum

software.intel.com/ai

Thank you

Learn more at ai.intel.com