INTEL'S MACHINE LEARNING STRATEGY

Gary Paek, HPC Marketing Manager, Intel Americas

HPC User Forum, Tucson, AZ

April 12, 2016

TAXONOMIC FOUNDATIONS

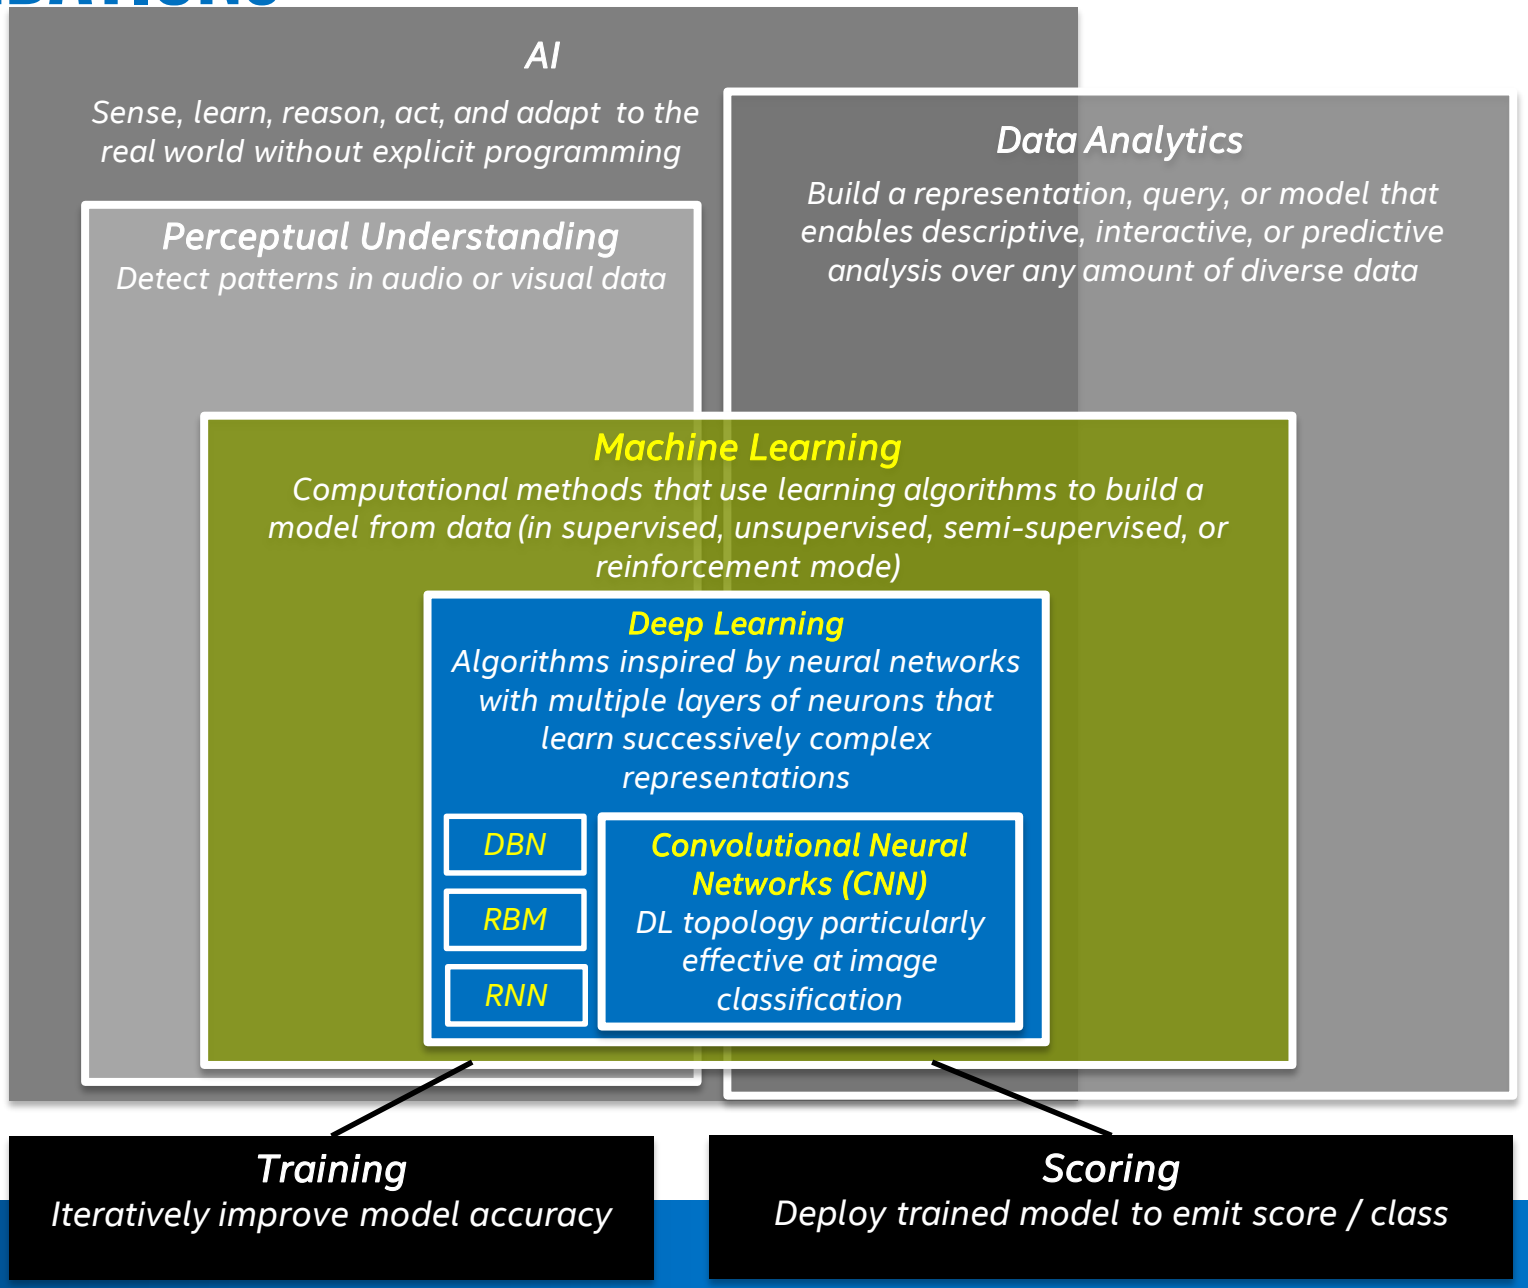

END TO END MACHINE LEARNING WORKFLOW

Model Scoring
Application SDK
Embedded OS

Model Deployment
Over-the-Air
Secure, Real Time

Model Performance
Tune for Performance
Longitudinal Analysis

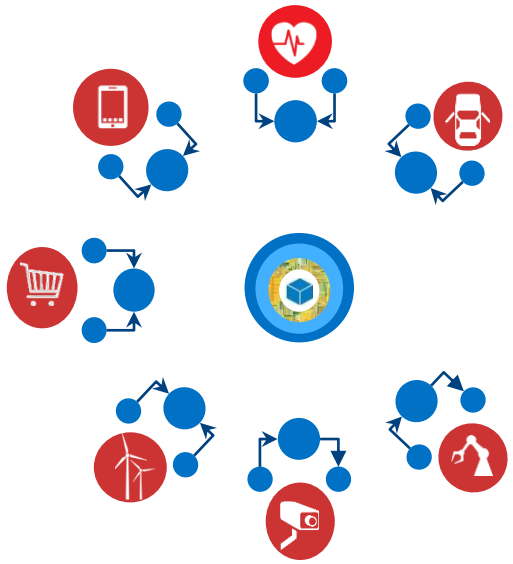

Things

New Data

Model Update
Track Model Drift
Manage Model Lifecycle

Model Development
Train for Accuracy

Data Annotation
Label Data

Data Curation
Inventory Data Sets

Data Acquisition

Data Aggregation

WHY NOW?

Bigger Data

Numbers: 5 KB / record
Text: 500 KB / record
Image: 1000 KB / picture
Audio: 5000 KB / song
Video: 5,000,000 KB / movie
High-Res: 50,000,000 KB / object

Better Hardware

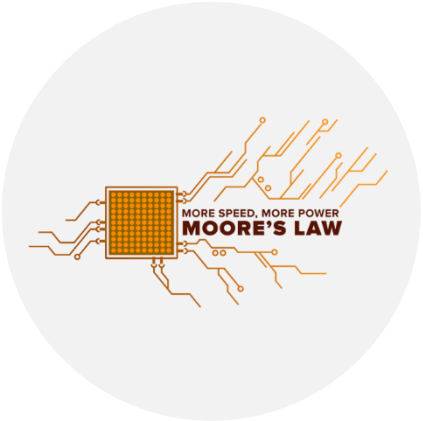

Transistor density doubles 18m
Computation / kwh doubles 18m
Cost / Gigabyte in 1995: $1000.00
Cost / Gigabyte in 2015: $0.03

Smarter Algorithms

Theoretical advances in training multi-layer feedforward neural networks led to better accuracy

New mathematical techniques for optimization over non-convex curves led to better learning algorithms

MACHINE LEARNING | APPLICABILITY

	Application	Model Type
IMAGENET	*Object Localization and Image Classification*	*Convolutional Neural Networks (CNN), Support Vector Machines*
NETFLIX RECOMMENDED TV	*Collaborative Filtering, Recommendation Engines, Inputting Missing Interactions*	*Restricted Boltzmann Machines (RBM), ALS*
	Anomaly Detection	*Clustering, Decision Trees*
"Ok Google"	*Forecasting or prediction of time-series and sequences like speech and video*	*Recurrent Neural Networks (RNN), Long-short Term Memory (LSTM), Hidden Markov Models*
CTR 	*Click Through Rate (CTR) Prediction*	*Logistic Regression*
	*State-Action Learning, Decision Making*	*Deep Q Networks (Reinforcement Learning)*

WHY INTEL?

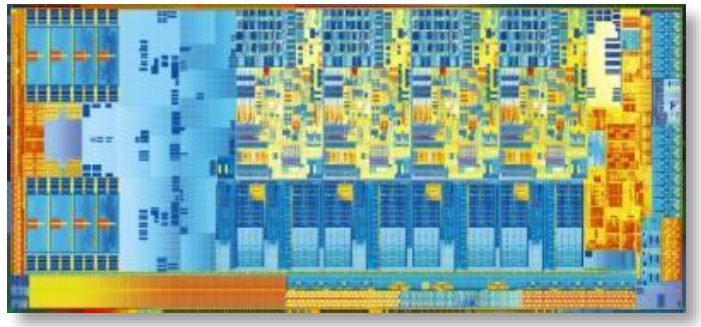

TRAINING **MACHINE LEARNING** INFERENCE

intel ITANIUM inside | intel XEON PHI inside | intel XEON inside | intel CORE i5 inside | intel ATOM inside | intel QUARK inside

INTEL MACHINE LEARNING STRATEGY

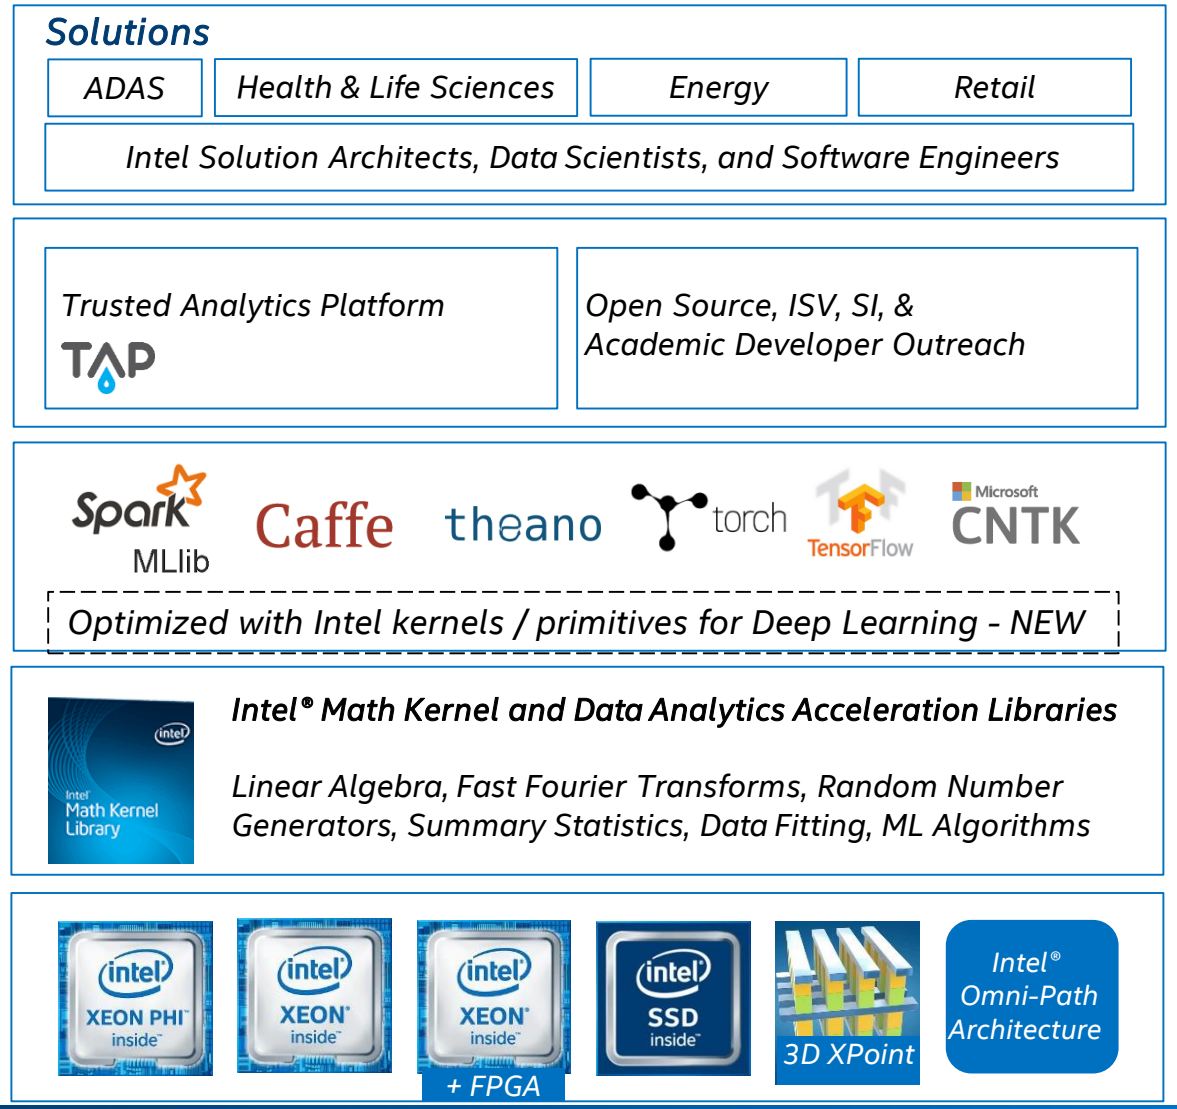

*Support the industry **innovation** across verticals*

***Accelerate** adoption by providing tools to the ecosystem*

*Enable and **optimize** key industry frameworks*

*Extract maximum **performance** through libraries*

*Enable optimization of **single-node and cluster** performance for Compute, Networking and Storage*

INTEL INVESTMENT IN HPC

LEADERSHIP

Intel® Omni-Path Fabric

Intel® Solutions for Lustre Software*

3D XPoint™ Technology

Intel® Parallel Studio

Intel® Cluster Studio

INNOVATIVE SOLUTIONS

HPC

High Performance Data Analytics

Single System Architecture

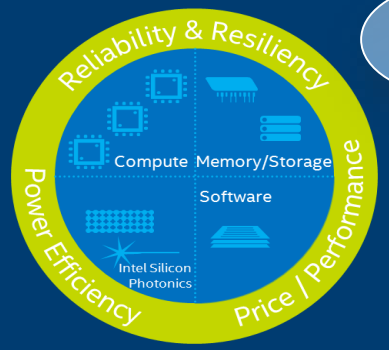

HPC Optimized

Extend Solution to Cloud

Intel® Scalable System Framework

SW ECOSYSTEM

Modernizing Community Code

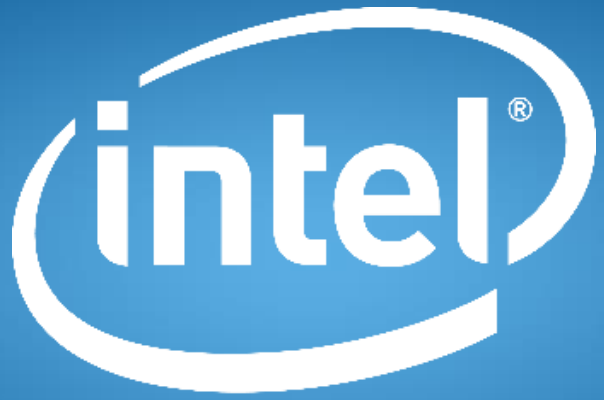
intel
experience
what's inside™

Machine/Deep Learning | Resources

Intel Caffe Repo: (Support for Multi-node Training)
https://github.com/intelcaffe/caffe

Spark MLLib Repo:
http://spark.apache.org/mllib/

Intel Machine Learning Blog Posts:
Myth Busted - CPUs and Neural Network Training
Caffe Scoring on Xeon Processors
Caffe Training on Multi-node Distributed Memory Systems

Trusted Analytics Platform:
http://trustedanalytics.org/

Performance Libraries:
MKL for Neural Networks - Technical Preview
Math Kernel Library
MKL Community License
Data Analytics Acceleration Library

Legal Disclaimers

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 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 {most relevant URL to your product}.

Intel, the Intel logo, {List the Intel trademarks in your document} 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.*