

ARTIFICIAL INTELLIGENCE AND HPC

Matt Beale Al Products Group June 2018

ARTIFICIAL INTELLIGENCE LANDSCAPE

DATA CENTER	EDGE	END POINT
Servers & appliances	Servers, appliances & gateways	IOT Sensors (Security, SDC Automotive
		Vision & Speech Vision as a Service
Dynamic, hybrid Dedicated DL	Streaming, latency and power	Desktop & Mobile
and analytics	and on premise	
systems		Display, video, AR/VR, rendering
DEEP I FARNING SCIENTIST	DATA SCIENTIST	SOFTWARE DEVELOPER
Values: Innovating models using the best tools and open sourced code	Values: Interactive tools for fast prototyping; rich set of connectors to different data sources	Values: Productivity & SW performance; right tools for platform; easy deployment and scaleImage: Control of the tools for platform; easy tools for platform; easy tools for platform; easy deployment and scale
Uses optimized frameworks such as Caffe TensorFlow, Theano, MxNet	Uses programming tools: R, Matlab, (NumPy, SciPy, scikit-learn), Keras	Python, Uses big data tools such as SPARK, Hadoop, and BigDL Framework



DEEP LEARNING DEPLOYED



Source: Intel customer engagement



(intel) Al 3





HPC & AI CONVERGENCE

Data analytics workflows becoming increasingly complex, with HPC-like modeling phases applied to large data sets

HPC increasingly leveraging Data Analytics to glean added insight into scientific and engineering data sets.



5

(intel)

HPC AND DEEP LEARNING



Statistical Machine Learning Create a catalog of every object in the visible universe from the Sloan Digital Sky Survey (SDSS)

Problem: What is a star? A galaxy? What overlaps?

Machine Learning: "Celeste" project uses Bayesian inference on pixel intensities to infer the object

Result: Created catalog in <15 minutes using a 55 terabyte data set on the NERSC Cori supercomputer



HPC and Deep Learning Complex Global Climate Model generates 400km view of atmospheric rivers and severe weather patterns

Deep Learning across 9600 nodes of the Cori supercomputer generates 25km view – and new insight

6

Learn more: https://arxiv.org/abs/1801.10277	Learn more: https://arxiv.org/abs/1708.05256
Cataloging the Visible Universe through Bayesian Inference at Petascale	Deep Learning at 15PF: Supervised and Semi-Supervised Classification for Scientific Data
Image courtesy of the National Energy Research Scientific Computing Center	Image courtesy of the National Energy Research Scientific Computing Center
mage counted, of the hadronal first g, herear of occurrence comparing conten	

Multi-Node Intel[®] Optimized Caffe ResNet-50 Scaling Efficiency on 2S Intel[®] Xeon[®] Platinum 8160 Processor Cluster MareNostrum4/BSC*



Source SURFsara blog: https://blog.surf.nl/en/imagenet-1k-training-on-intel-xeon-phi-in-less-than-40-minutes/

*MareNostrum4 (Barcelona Supercomputing Center): https://www.bsc.es/marenostrum

Configuration Details on Slide: MareNostrum4/BSC* Configuration Details:

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. Copyright © 2017, Intel Corporation

