

SUPERCOMPUTING IMPACT IN BIG DATA & ARTIFICIAL INTELLIGENCE

PROF. DR. – ING. MORRIS RIEDEL, UNIVERSITY OF ICELAND / JUELICH SUPERCOMPUTING CENTRE HEAD OF CROSS-SECTIONAL TEAM DEEP LEARNING & RDA CO-CHAIR INTEREST GROUP BIG DATA 25TH JULY, KEYNOTE, NEXTGEN @ HELMHOLTZ 2018 CONFERENCE, POTSDAM, GERMANY













JÜLICH SUPERCOMPUTING CENTRE

FORSCHUNGSZENTRUM JUELICH (FZJ)

Multi-Disciplinary Research Centre of the Helmholtz Association in Germany



(Juelich Supercomputing Centre known as JSC)

- Selected Facts
 - One of EU largest inter-disciplinary research centres (~5000 employees)



 Expertise in physics, materials science, nanotechnology, neuroscience and medicine & information technology (HPC, Big Data, Artificial Intelligence) [1] Holmholtz Association Web Page

25th July 2018 Page 2

HIGH PERFORMANCE COMPUTING (HPC)

Perspective: Floating Point Operations per one second (FLOPS or FLOP/s)

1.000.000 FLOP/s

- e Photograph by Rama, Wikimedia Commons
- 1 GigaFlop/s = 10⁹ FLOPS
- 1 TeraFlop/s = 10¹² FLOPS
- 1 PetaFlop/s = 10^{15} FLOPS
- 1 ExaFlop/s = 10^{18} FLOPS

1.000.000.000.000.000 FLOP/s → 1 PFLOP/s

~295.000 cores~2009 (JUGENE)



25th July 2018



TOP500 SUPERCOMPUTERS – JUNE 2018

Enabling High Performance Computing

Rank	System		Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)	
1	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States			122,300.0	187,659.3	8,806	
2	Sunway TaihuLight - Sunway MPP, Sur 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wu China	nway SW26010 260C uxi	10,649,600	93,014.6	125,435.9	15,371	
3	Sierra - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/NNSA/LLNL United States		1,572,480	71,610.0	119,193.6		
23 [27]	Forschungszentrum Juelich (FZJ) Germany Top500	JUWELS Module 1 - Bull Sequana X1000, Xeon Plat 8168 24C 2.7GHz, Mellano InfiniBand/ParTec ParaSta ClusterSuite Bull, Atos Group	114,4 inum x EDR ation	80 6,177.5	7 9,891.1	1,361	
	25 th July 2018 Page 4						





[9] Distributed & Cloud Computing Book

- Top500 #1 Summit (ORNL) 6 GPUs/node (June 2018)
- 1st time more flop/s added by GPUs vs. CPUs (2018)
- The LINPACK performance benchmark not fully reflects the broad range of applications in HPC today

SIMULATION SCIENCES

Traditional Supercomputing Impact in Scientific Computing



- Known physical laws
- Numerical methods
- Parallel Computing





RESEARCH FIELDS AND THEIR SHARES (2016-2018)

Supercomputing Systems Utilization → Requirement for sending in HPC project proposals to receive 'time'



HPC DRIVES 'BIG DATA' STORAGE





- Scientific computing applications can be considered as'creators' of big data in science & engineering domains
 Traditional simulation
 - Traditional simulation sciences already require high capacity in data storages to store output of various application models (or checkpointing)

'BIG DATA' INFRASTRUCTURE FOR HPC & DATA SCIENCE

Multiple systems combined with whole federation of other Helmholtz centre systems

JUST Storage Cluster

- IBM Spectrum Scale file system (GPFS)
- 75 PB gross capacity
- 5th generation
- Parallel access
- Tape Libraries
 - Automated cartridge systems
 - 300 PB
 - 3 libraries (in 2 buildings)
 - 60 tape drives
 - 35,000 tapes



- Journals will/do require to store data with persistent identifiers for papers
- Need for sharing of data in communities (e.g. turbulence flow)



25th July 2018

EUROPEAN UNION & COMMISSION PLANS

Strategic Plans towards Artificial Intelligence & Supercomputers

"By supporting strategic projects in frontline areas such as artificial intelligence, supercomputers,

cybersecurity or industrial digitisation, and investing in digital skills, the new programme will help to complete the Digital Single Market, a key priority of the Union."

[2] COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT. THE EUROPEAN COUNCIL, THE COUNCIL, THE EUROPEAN ECONOMIC AND SOCIAL COMMITTEE AND THE COMMITTEE OF THE REGIONS. EC, 2018, 2nd May 2018

13 You Retweeted

European Commission 🔲 🥺 @EU Commission · Apr 25 How can Europe be at the forefront of artificial intelligence, #dataeconomy and digital healthcare?

@Ansip EU @GabrielMariya @EBienkowskaEU @Moedas #DigitalSingleMarket #AI



Digital Single Market proposals: artificial intelligence, data econ... European Commission @EU Commission





Follow

We are proud of you @fzi isc for the #firstclass #supercomputing facility you run. It is by efforts like yours that we reaffirm #EUaddedvalue and leadership in groundbreaking technologies. It is by cooperating that we will achieve our objectives for #EU leader in #HPC



8:28 AM - 5 Mar 2018

25th July 2018

ARTIFICIAL INTELLIGENCE & DATA SCIENCE

- 1. Some pattern exists
- 2. No exact mathematical formula
- 3. Data exists
- Idea 'Learning from (Big) Data' shared with a wide variety of other disciplines
 - E.g. signal processing, data mining, etc.
 - Challenges: too slow, not enough memory, etc.
 - Artificial Intelligence uses methods from machine learning, pattern recognition, data mining, applied statistics, deep learning, etc.

Page 10

 Artificial Intelligence is a very broad subject and goes from very abstract theory to extreme practice ('rules of thumb')



INNOVATIVE DEEP LEARNING TECHNOLOGIES

Recent Innovative & Disruptive Approach in Artificial Intelligence



Engeneer

Common Deep Learning frameworks like Keras/TensorFlow take advantage of HPC via GPUs

> Workshops tomorrow cover deep learning technologies and approaches as well as machine learning methods

25th July 2018

RECENT TREND IN SCIENCE: PYTHON

SIMPLE AND FLEXIBLE PROGRAMMING LANGUAGE

- Selected Benefits
 - Work with many students reveal: qucky & easy to learn
 - Is an interpreted powerful programming language
 - Has Efficient high-level data structures
 - Provides a simple but effective approach to object-oriented programming
 - Great libraries & community support (e.g. numpy)
 - Python is an ideal language for fast scripting and rapid application development that in turn makes it interesting for the machine learning modeling process
 - The machine learning modeling process in general and the deep learning modeling process in particular requires iterative and highly flexible approaches
 - E.g. network topology prototyping, hyper-parameter tuning, etc.



outhon





[6] F. Chollet, 'Deep Learning with Python' Book

DEEP LEARNING FRAMEWORKS

Example: Programming with TensorFlow & Python

- Tensorflow is an open source library for deep learning models using a flow graph approach
- Tensorflow nodes model mathematical operations and graph edges between the nodes are so-called tensors (also known as multi-dimensional arrays)
- The Tensorflow tool supports the use of CPUs and GPUs (much more faster than CPU versions)
- Tensorflow work with the high-level deep learning tool Keras in order to create models fast (i.e. layer-wise fashion)



DEEP LEARNING

Learning process using massive amounts of Matrix/Vector – Matrix multiplications

- A Tensor is nothing else than a multi-dimensional array often used in scientific & engineering environments
- Tensors are best understood when comparing it with vectors or matrices and their dimensions
- Those tensors 'flow' through the deep learning network during the optimization / learning & inference process





(matrix of dimensions [5,6])





(three dimensional tensor) (tensor of dimension [4,4,3])



[1] M. Riedel, Invited YouTube Tutorial on Deep Learning, Ghent University

25th July 2018

DEEP LEARNING

Programming with TensorFlow & Keras

- Keras is a high-level deep learning library implemented in Python that works on top of existing other rather lowlevel deep learning frameworks like Tensorflow, CNTK, or Theano
- The key idea behind the Keras tool is to enable faster experimentation with deep networks
- Created deep learning models run seamlessly on CPU and GPU via low-level frameworks

keras.layers.Dense(units,

activation=None, use_bias=True, kernel_initializer='glorot_uniform', bias_initializer='zeros', kernel_regularizer=None, bias_regularizer=None, activity_regularizer=None, kernel_constraint=None, bias_constraint=None)

keras.optimizers.SGD(lr=0.01,

momentum=0.0,
decay=0.0,
nesterov=False)

- Tool Keras supports inherently the creation of artificial neural networks using Dense layers and optimizers (e.g. SGD)
- Includes regularization (e.g. weight decay) or momentum

[5] M. Riedel, Invited YouTube Tutorial on Deep Learning, Ghent University



[10] Keras Python Deep Learning Library

25th July 2018

DEEP LEARNING – ARCHITECTURES

Each of the Architectures provide Unique Characteristica (e.g. 'smart layers')

- Deep Neural Network (DNN)
 - 'Shallow ANN' approach with many hidden layers between input/output
- Convolutional Neural Network (CNN, sometimes ConvNet)
 - Connectivity pattern between neurons is like animal visual cortex





- Deep Belief Network (DBN)
 - Composed of mult iple layers of variables; only connections between layers
- Recurrent Neural Network (RNN)
 - 'ANN' but connections form a directed cycle; state and temporal behaviour
 25th July 2018 Page 16
- Deep Learning architectures can be classified into Deep Neural Networks, Convolutional Neural Networks, Deep Belief Networks, and Recurrent Neural Networks all with unique characteristica
- Deep Learning needs 'big data' to work well & for high accuracy – works not well on sparse data

SUPERCOMPUTING IMPACT VIA GPUS

Disruptive results for Deep Learning – Now State-of-the-Art

- Dataset: ImageNet
 - Total number of images: 14.197.122
 - Number of images with bounding box annotations: 1.034.908

Number synsets in	the subtree).	Treemap Visualization Ima	ges of the Synset	Downloads	
+ Imag	eNet 2011 Fall Release (32326) lant, flora, plant life (4486) eological formation, formation (1 atural object (1112) port, athletics (176)		U		
*- a	trtract, artefact (10504) - instrumentality, instrumentation - device (2760) - musical instrument, inst - acoustic device (27)				
			•т	raditional CV 🛛 🛑 Deep I	Learning
	79% •				-
	70%				
	60% 🎖				
Rate	50%				
Error	40%				
	30%			ę	
	20%			8	
	10% 7%			8	
	2010	2011	2012	2013	2014



High level category	# synset (subcategories)	Avg # images per synset	Total # images
amphibian	94	591	56K
animal	3822	732	2799K
appliance	51	1164	59K
bird	856	949	812K
covering	946	819	774K
device	2385	675	1610K
fabric	262	690	181K
fish	566	494	280K
flower	462	735	339К
food	1495	670	1001K
fruit	309	607	188K
fungus	303	453	137K
furniture	187	1043	195K
geological formation	151	838	127K
invertebrate	728	573	417K
mammal	1138	821	934K
musical instrument	157	891	140K
plant	1666	600	999K
reptile	268	707	190K
sport	166	1207	200K
structure	1239	763	946K
tool	316	551	174K
tree	993	568	564K
utensil	86	912	78K
vegetable	176	764	135K
vehicle	481	778	374K
person	2035	468	952K

[23] ImageNet Web page

DEEP LEARNING & GPU PARALLELIZATION

speedup

1.0

3.2

6.5

12.5

22.4

41.4

101.1

images/s

55

178

357

689

1230

2276

5562

Simple Image Benchmark on JURECA JSC HPC System (75 x 2 NVIDIA Tesla K80/node – dual GPU design)

Performance per GPU [images/s]

second and relative speedup

55

44.5

44.63

43.06

38.44

35.56

43.45

- Setup
 - TensorFlow 1.4
 - Python 2.7
 - CUDA 8
 - cuDNN 6
 - Horovod 0.11.2
- [21] A. Sergeev, M. Del Balso,'Horovod', 2018

#GPUs

1

4 8

16

32

64

128

- MVAPICH-2.2-GDR
- 'Simple' 1.2 million images with 224 x 224 pixels
- Tool Horovod (using Message Passing Interface) for distributed deep learning TensorFlow (and Keras)
- Machine & Deep Learning: speed-up is just secondary goal after primary goal accuracy (applications!)
- Speed-up & parallelization nice to have for faster hyperparameter tuning, model creation, and inference
- Third goal is avoiding much feature engineering through 'feature learning' of deep learning networks

Nvidia NVLink will be next interesting technology



[20] JURECA HPC System

COMPARE TRADITIONAL MACHINE LEARNING

Supervised Classification Example – Remote Sensing Dataset & Results

- Traditional Methods
 - Support Vector Machine (SVM)
 - 52 classes of different land cover, 6 discarded, mixed pixels, rare groundtruth

Data

Engeneer Transform

Reduce

Learn

Features

Model

- Substantial manual feature engineering, e.g. Self Dual Attribute Profile (SDAP)
- 10-fold cross-validation
- Achieved 77,02 % accuracy





Traditional

Machine

Learning

[15] G. Cavallaro and M. Riedel, et al. , 2015



SUPERCOMPUTING IMPACT IN MACHINE LEARNING

Engeneer

Supervised Classification Example – Speed-up Cross-Validation & MPI C Code

- **Traditional Methods**
 - Support Vector Machine (SVM)
 - pISVM C code that can be improved (taken from ML experts - not parallel experts, tuned @ JSC)
 - Message Passing Interface (MPI)
- Feature Engineering
 - Riedel, et al., 2018 Working also on parallel feature engineering using tree-based approach (MPI/OpenMP C)

0000



 C_{3}^{240}

 C_5^{174} C_5^{127}

 C_{4}^{212} C_{5}^{0}



[16] M. Goetz and M.

Scenario 'pre-processed data', 10xCV serial: accuracy (min)

γ/C	1	10	100	1000	10 000
2	48.90 (18.81)	65.01 (19.57)	73.21 (20.11)	75.55 (22.53)	74.42 (21.21)
4	57.53 (16.82)	70.74 (13.94)	75.94 (13.53)	76.04 (14.04)	74.06 (15.55)
8	64.18 (18.30)	74.45 (15.04)	77.00 (14.41)	75.78 (14.65)	74.58 (14.92)
16	68.37 (23.21)	76.20 (21.88)	76.51 (20.69)	75.32 (19.60)	74.72 (19.66)
32	70.17 (34.45)	75.48 (34.76)	74.88 (34.05)	74.08 (34.03)	73.84 (38.78)

Scenario 'pre-processed data', 10xCV parallel: accuracy (min)

γ /C	1	10	100	1000	10 000
2	75.26 (1.02)	65.12 (1.03)	73.18 (1.33)	75.76 (2.35)	74.53 (4.40)
4	57.60 (1.03)	70.88 (1.02)	75.87 (1.03)	76.01 (1.33)	74.06 (2.35)
8	64.17 (1.02)	74.52 (1.03)	77.02 (1.02)	75.79 (1.04)	74.42 (1.34)
16	68.57 (1.33)	76.07 (1.33)	76.40 (1.34)	75.26 (1.05)	74.53 (1.34)
32	70.21 (1.33)	75.38 (1.34)	74.69 (1.34)	73.91 (1.47)	73.73 (1.33)

First Result: best parameter set from 14.41 min to 1.02 min Second Result: all parameter sets from ~9 hours to ~35 min

[15] G. Cavallaro and M. Riedel, et al., 2015

TRADITIONAL MACHINE LEARNING VS DEEP LEARNING

Supervised Classification Example – Remote Sensing Dataset & Results

- Traditional Methods
 - C MPI-based Support Vector Machine (SVM)
 - Substantial manual feature engineering
 - 10-fold cross-validation for model selection
 - Achieved 77,02 % accuracy
- Convolutional Neural Networks (CNNs)
 - Python/TensorFlow/Keras
 - Automated feature learning
 - Achieved 84,40 % accuracy on all 58 classes



Hyperspectral Image Cube



3 x

Engeneer

[12] J. Lange, G. Cavallaro, M. Riedel, et al. , 2018

DEEP LEARNING HYPERPARAMETER TUNING

'Finding' good Deep Learning Network Topology requires Supercomputing



- Using Python with TensorFlow & Keras easily enables changes in hyper-parameter tuning
- Various runs on different topologies add up to computational demand of (interlinked) GPUs
- Need for HPC machines with good GPUs and good deep learning software stacks required



SUPERCOMPUTING IMPACTS IN ARTIFICIAL INTELLIGENCE

Modular Supercomputing Design for Resource Provisioning in Helmholtz Association



MODULAR SUPERCOMPUTING ARCHITECTURE



DEEP

25th July 2018





 Example: Weather Research & Forecasting (WRF) models

25th July 2018

Page 27

[3] Thomas Lippert, Daniel Mallmann, Morris Riedel Publication Series of the John von Neumann Institute for Computing (NIC) NIC Series 48, 417, ISBN 978-3-95806-109-5, pp. 1 – 10, 2016

OTHER INTERESTING POINTERS

Helmholtz Association Activities & Research Data Alliance

- Helmholtz Data Federation (HDF)
 - Federation and extension of multi-topical data centers with new storage- and analysis hardware
 - Usage of innovative data management solutions & excellent user support
- Helmholtz Analytics Framework (HAF)
 - Common components for data analytics
 - Applied parallel machine learning methods
- Research Data Alliance (RDA)
 - Research data sharing without boundaries
 - Interest groups and working groups

ochastik & Statistics

pervised learning

supervised learning

timisatio

age Analysi



SUMMARY

Mindset

- Think traditional machine learning still relevant for deep learning & big data analysis
- Using interpreted languages like Python enable easy & flexible parameter-tuning
- Selected new approaches with specific deep learning per problem (CNN, LSTM, etc.)

Skillset

- Basic knowledge of machine learning required for deep learning
- Faster experimentation and use of various topologies and architectures through Python
- Validation (i.e. model selection) and regularization still valid(!)
- Toolset
 - Parallel versions of traditional machine learning methods exist
 - Python with Tensorflow & Keras just one example that takes advantage of supercomputing
 - Explore technology trends, e.g. specific chips for deep learning, NAM, NVLink, etc. 25th July 2018 Page 29









REFERENCES



REFERENCES (1)

- [1] Helmholtz Association Web Page, Online: https://www.helmholtz.de/en/
- [2] European Commission, COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT, THE EUROPEAN COUNCIL, THE COUNCIL, THE EUROPEAN ECONOMIC AND SOCIAL COMMITTEE AND THE COMMITTEE OF THE REGIONS, EC, 2018, 2nd May
 Online: https://ec.europa.eu/commission/sites/beta-political/files/communication-modern-budget-may2018 en.pdf?utm source=POLITICO.EU&utm campaign=e3a8a86cc6 EMAIL CAMPAIGN 2018 05 02&utm medium=email&utm term=0 10959edeb5-e3a8a86cc6-189710085
- [3] Thomas Lippert, Daniel Mallmann, Morris Riedel
 Publication Series of the John von Neumann Institute for Computing (NIC) NIC Series 48, 417, ISBN 978-3-95806-109-5, pp. 1 10, 2016
- [4] DEEP Projects Web Page,
 Online: <u>http://www.deep-projects.eu/</u>
- [5] M. Riedel, 'Deep Learning using a Convolutional Neural Network', Ghent University, Invited YouTube Tutorial, Online: <u>https://www.youtube.com/watch?v=gOL1_YIosYk&list=PLrmNhuZo9sgZUdaZ-f6OHK2yFW1kTS2qF</u>
- [6] François Chollet 'Deep Learning with Python', Book, ISBN 9781617294433, 384 pages, 2017, Online: <u>https://www.manning.com/books/deep-learning-with-python</u>
- [7] Tensorflow Deep Learning Framework, Online: <u>https://www.tensorflow.org/</u>

25th July 2018 Page 31

REFERENCES (2)

- [8] A Tour of Tensorflow,
 Online: https://arxiv.org/pdf/1610.01178.pdf
- [9] K. Hwang, G. C. Fox, J. J. Dongarra, 'Distributed and Cloud Computing', Book, Online: <u>http://store.elsevier.com/product.jsp?locale=en_EU&isbn=9780128002049</u>
- [10] Big Data Tips, 'What is a Tensor?',
 Online: <u>http://www.big-data.tips/what-is-a-tensor</u>
- [11] Keras Python Deep Learning Library, Online: <u>https://keras.io/</u>
- [12] J. Lange, G. Cavallaro, M. Goetz, E. Erlingsson, M. Riedel, 'The Influence of Sampling Methods on Pixel-Wise Hyperspectral Image Classification with 3D Convolutional Neural Networks', Proceedings of the IGARSS 2018 Conference, to appear
- [13] G. Cavallaro, N. Falco, M. Dalla Mura and J. A. Benediktsson, "Automatic Attribute Profiles," in IEEE Transactions on Image Processing, vol. 26, no. 4, pp. 1859-1872, April 2017, Online: <u>http://ieeexplore.ieee.org/document/7842555/</u>
- [14] C. Cortes and V. Vapnik, "Support-vector networks," Machine Learning, vol. 20(3), pp. 273–297, 1995
- [15] G. Cavallaro, M. Riedel, J.A. Benediktsson et al., 'On Understanding Big Data Impacts in Remotely Sensed Image Classification using Support Vector Machine Methods', IEEE Journal of Selected Topics in Applied Earth Observation and Remote Sensing, 2015, DOI: <u>10.1109/JSTARS.2015.2458855</u>
- [16] M. Goetz, G. Cavallaro, T. Geraud, M. Book, M. Riedel, 'Parallel Computation of Component Trees on Distributed Memory Machines', Journal of Transactions on Parallel and Distributed Systems, 2018, to appear

25th July 2018 Page 32

REFERENCES (3)

- [19] M. Goetz, G. Cavallaro, T. Geraud, M. Book, M. Riedel, 'Parallel Computation of Component Trees on Distributed Memory Machines', Journal of Transactions on Parallel and Distributed Systems, 2018, to appear
- [20] JURECA HPC System @ Juelich Supercomputing Centre (JSC),
 Online: http://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/JURECA/Configuration/Configuration_node.html;jsessionid=CF5F5276AE7CFF0B185DCA048A2D0010
- [21] A. Sergeev, M. Del Balso, 'Horovod: fast and easy distributed deep learning in TensorFlow', 2018 Online: <u>https://arxiv.org/abs/1802.05799</u>
- **[22]** J. Dean et al., 'Large scale deep learning', Keynote GPU Technical Conference, 2015
- [23] ImageNet Web page,
 Online: <u>http://image-net.org</u>
- [24] Top500 Supercomputers,
 Online: <u>https://www.top500.org</u>

25th July 2018

ACKNOWLEDGEMENTS

Previous & current members of the High Productivity Data Processing Research Group





PD Dr. G. Cavallaro



Senior PhD Senior PhD Student A.S. Memon Student M.S. Memon





PhD Student E. Erlingsson



PhD Student S. Bakarat



MSc Student G.S. Guðmundsson (Landsverkjun)



Dr. M. Goetz (now KIT)



MSc M. Richerzhagen



MSc P. Glock (now INM-1)



MSc C. Bodenstein (now Soccerwatch.tv)



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 763558

Further acknowledgements of this talk go to Prof. Dr. Dr. Thomas Lippert & the members of DEEP-EST EU Project

25th July 2018

THANKS

Talk shortly available under www.morrisriedel.de



25th July 2018