POLITECNICO DI MILANO

Department of Electronics, Informatics and Bioengineering

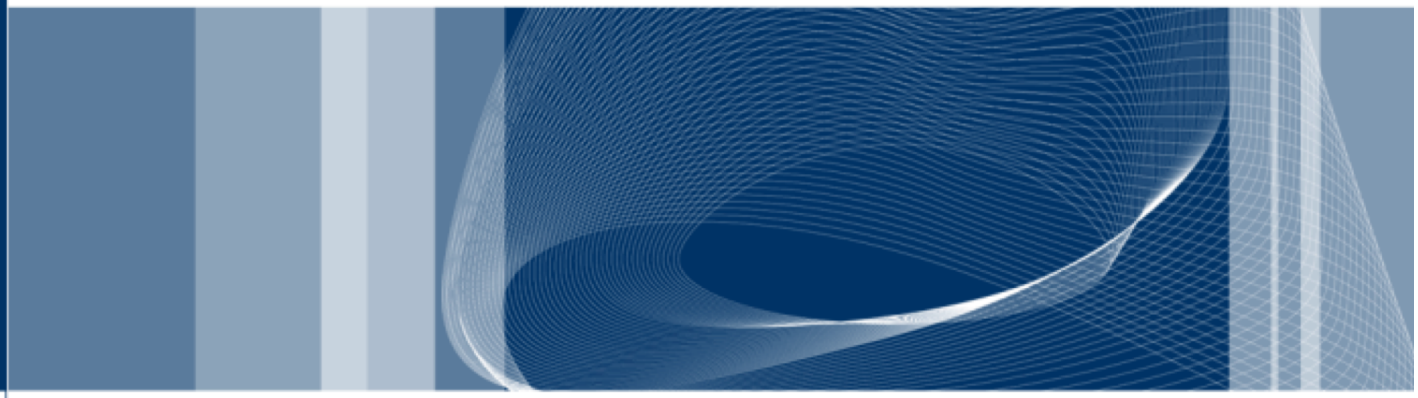

Machine learning for studying the performance of large scale systems

Danilo Ardagna, Marco Lattauda, Eugenio Gianniti

Why Machine Learning Models

- Big data applications, e.g., base on Apache Spark framework, are rapidly growing
  - Hundreds/thousands of stages
- Growing interest also on Deep Learning, e.g., Convolutional Neural Networks
  - Training on GPGPUs, thousands of cores in a single computational node
- Scalability issues of traditional analytical models
- Current focus:
  - Develop a benchmarking suite to automate the generation of performance profiles for ML training
  - Develop a library to automate the training of ML models and their hyper-parameter tuning

Sparkbench Library Extension

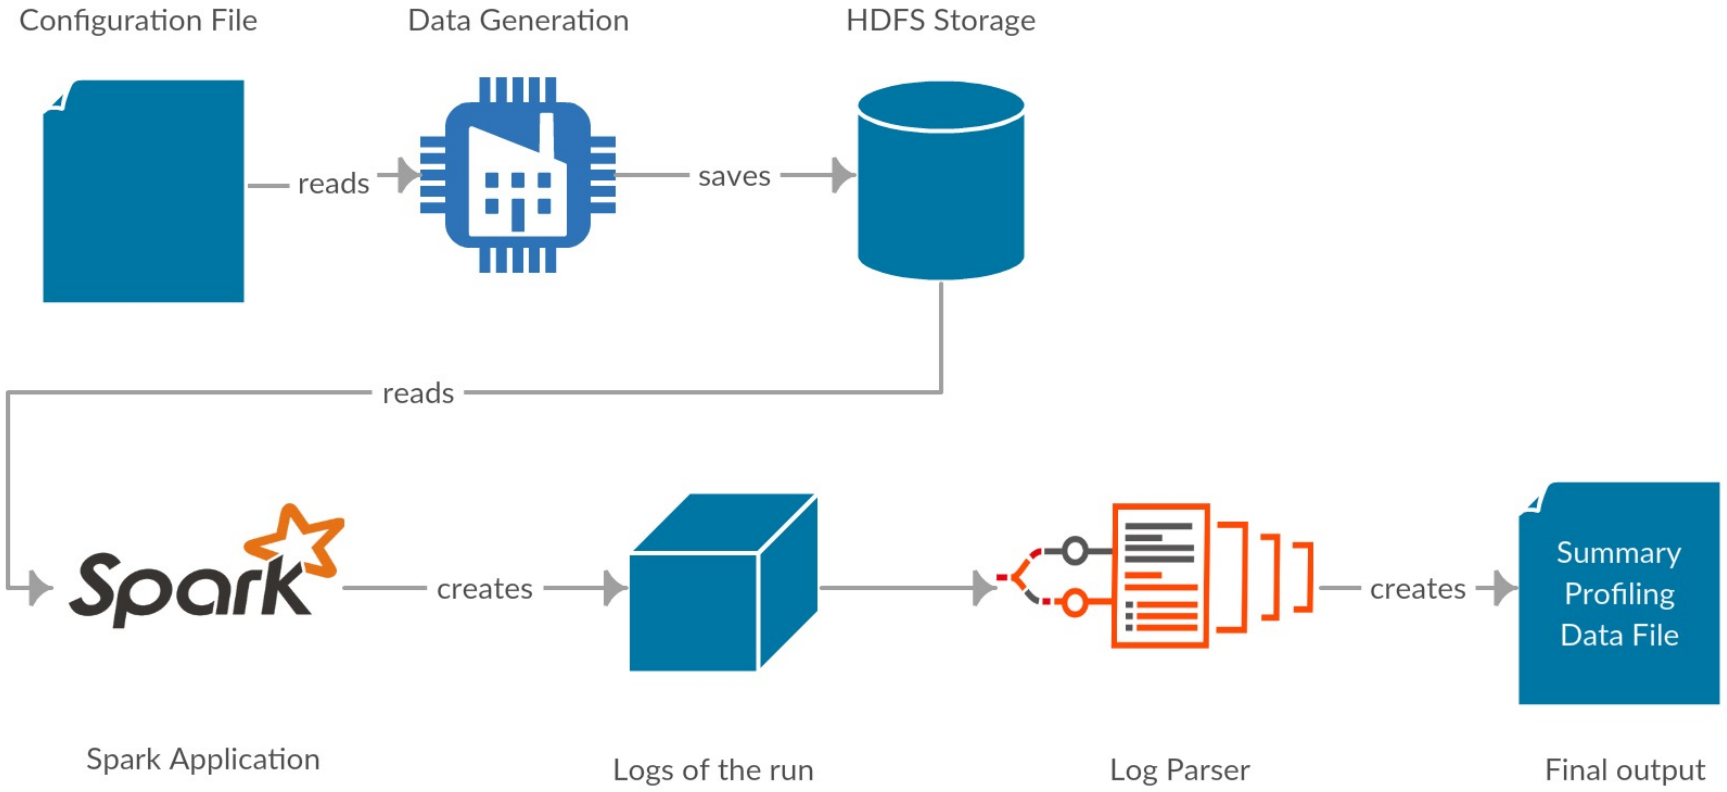

	Data Generation	Algorithm
Logistic Regression	✓	✓
Linear Regression	✓	✓
Random Forest	✓	✓
Spark DL		✓

Similar toolchain based on TPC-DS & CNNs

Machine Learning Model Library

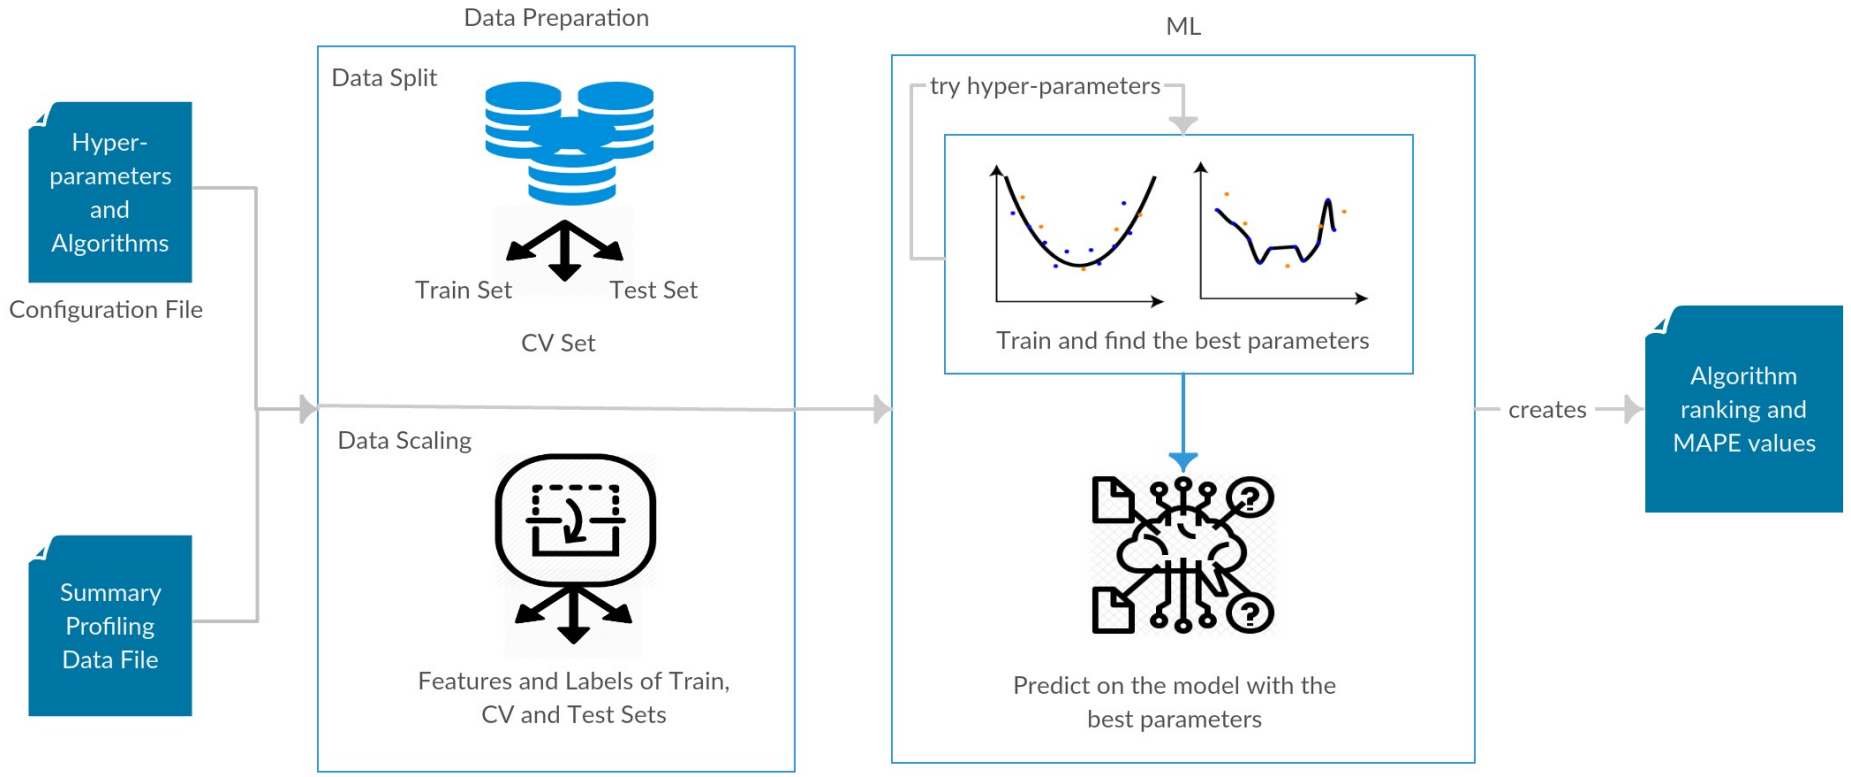

Experiment Settings

- **Workloads:**
  - TPC-DS - the industry benchmark for data warehouse systems
  - Sparkbench library
  - CNNs training on TensorFlow and Pytorch
- **Platforms:**
  - Microsoft HDInsigh
  - IBM Power8 Cluste
- **Evaluation Metrics:**
  - Mean Absolute Perc
  - Fraction of Large Er
- **Experiments:**
  - Interpolation and Ex
- **Regression Models:**

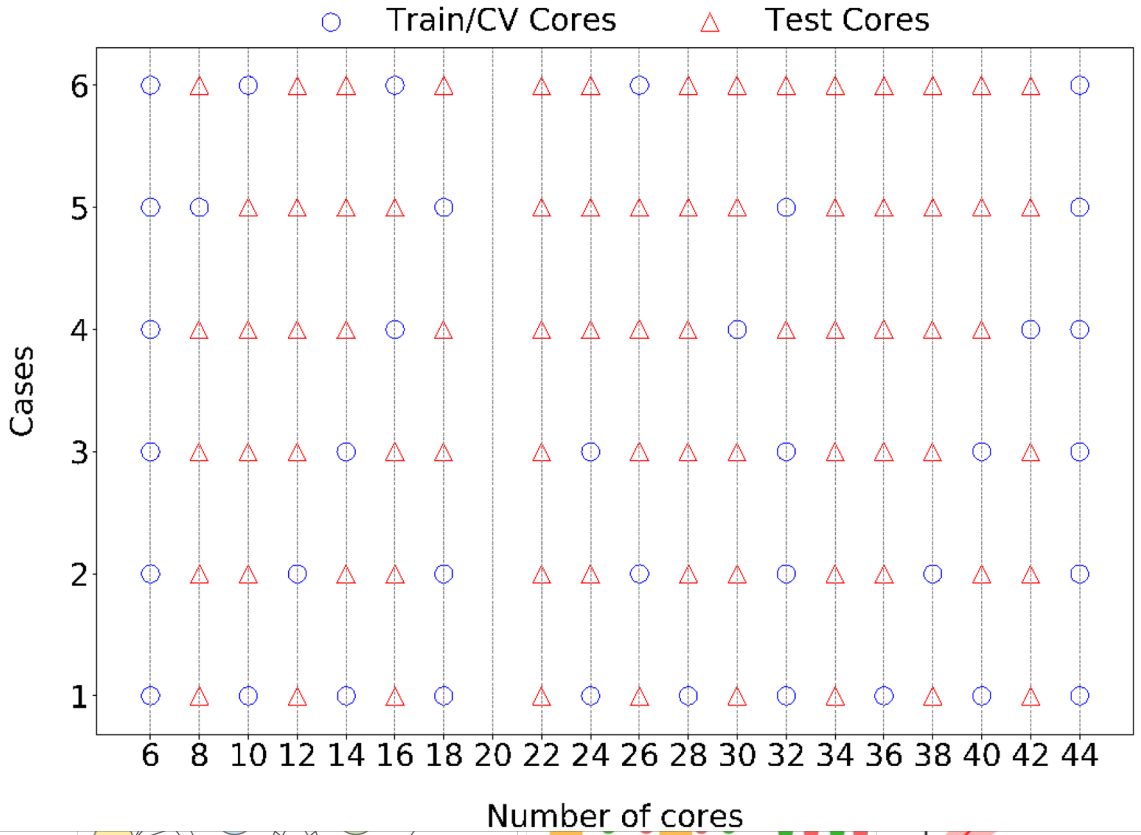

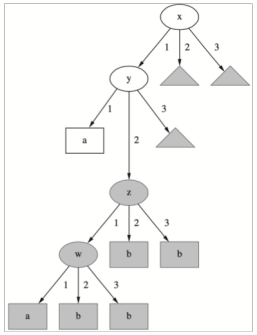

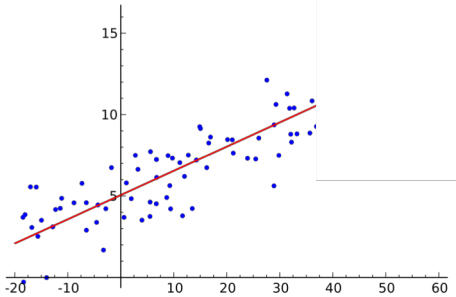

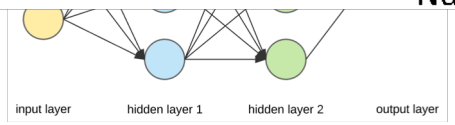

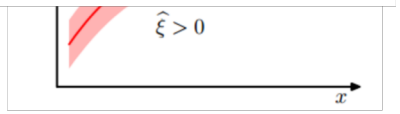

ML results - Spark

TPC-DS

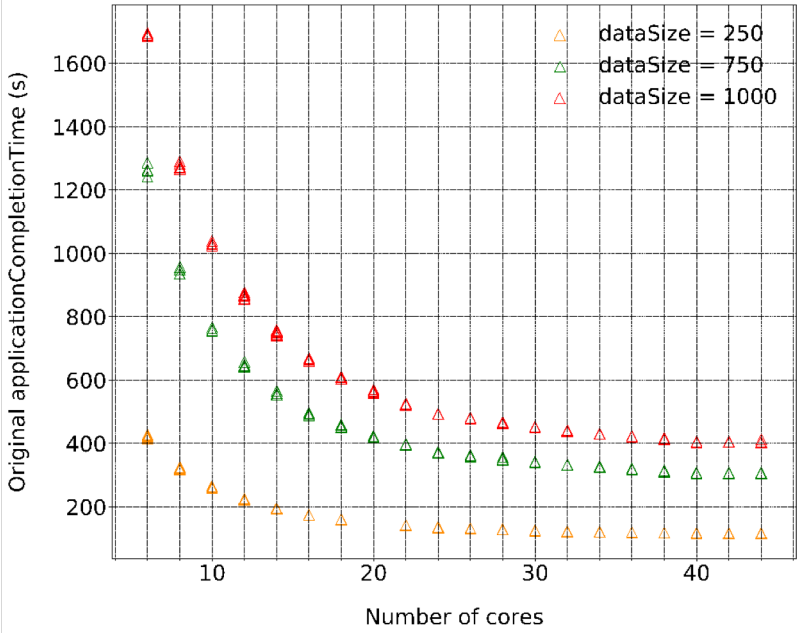

	Data Size in Train/CV Set	Data Size in Test Set
1	250 GB	250 GB
2	750 GB	750 GB
3	1000 GB	1000 GB
4	250 and 750 GB	1000 GB
5	250 and 1000 GB	750 GB
6	750 and 1000 GB	250 GB

@1 TB Conf 3

	Decision Tree	Linear Regression	Neural Network	Random Forest
Case 0	0.70	2.64	1.18	8.33
Case 1	8.62	5.69	3.36	11.26
Case 2	11.95	6.82	6.77	17.12
Case 3	15.41	5.64	6.48	16.48
Case 4	17.36	8.33	44.58	15.06
Case 5	24.70	33.58	11.92	33.17
Case 6	10.79	41.84	5.43	22.88

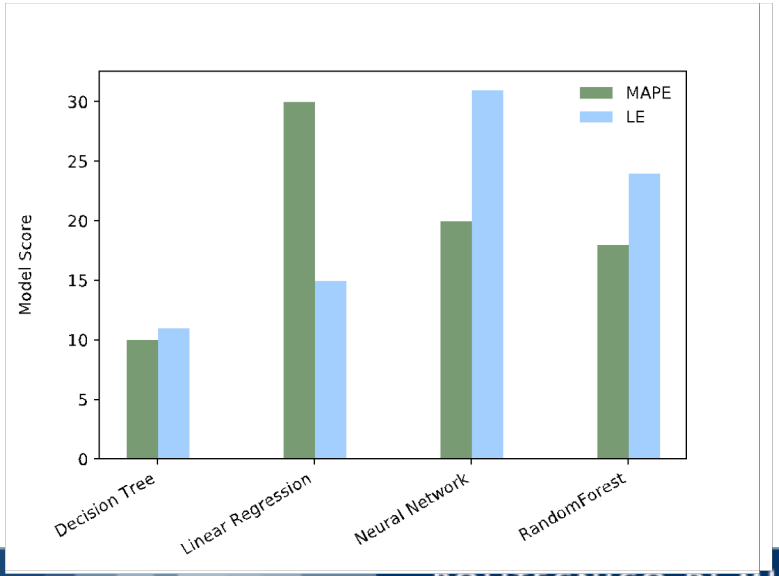

ML results - GPGPUs

Extrapolation goals:

- Change the batch size
- Exploitation of new hardware
- Training of new versions of a CNN

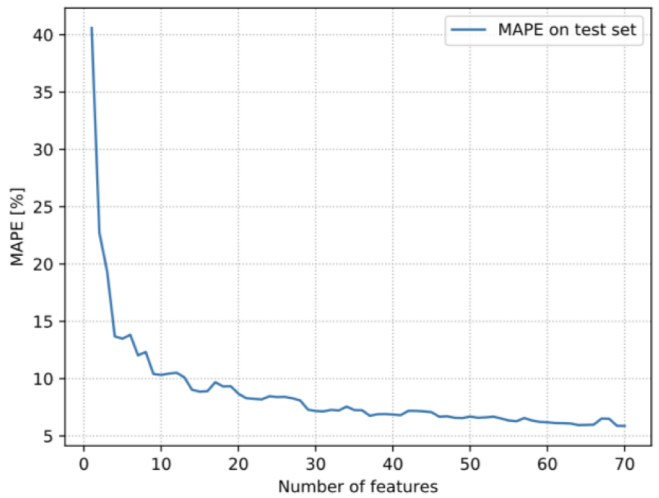

Network	Framework	MAPE
AlexNet	PyTorch	8.28
	TensorFlow	5.08
ResNet-50	PyTorch	18.09
	TensorFlow	10.10

GPUs number Extr.

Network	Framework	GPU Type K80	M60	GTX 1080Ti
AlexNet	PyTorch	7.21	12.18	4.98
	TensorFlow	24.75	17.27	8.77
ResNet-50	PyTorch	5.11	9.04	11.76
	TensorFlow	24.58	18.29	6.54
VGG-19	PyTorch	12.20	15.98	24.13
	TensorFlow	8.84	13.52	13.65

Extr. Inner Modules number

Network	Framework	Max N. IMs	GPU Type M60 1	2	4
ResNet	PyTorch	4	23.51	27.95	17.40
ResNet	PyTorch	5	24.85	25.11	16.75
ResNet	PyTorch	6	26.76	20.40	16.63
ResNet	PyTorch	8	17.06	7.93	15.99

Batch Size Extr.

Network	Framework	GPU Type P600 1	2	K80 1	2	3	4	M60 1	2	3	4	GTX 1080Ti 1	2	4	8
AlexNet	PyTorch	11.12	7.85	1.74	3.33	1.81	0.66	6.19	3.49	6.58	0.75	0.43	1.62	1.15	4.16
	TensorFlow	9.83	10.04	2.30	2.61	4.28	2.82	7.19	6.36	6.91	6.96	4.06	5.36	1.14	1.12
ResNet-50	PyTorch	10.64	11.97	0.76	7.83	3.09	4.53	3.60	20.04	9.58	4.64	12.62	11.93	20.63	4.29
	TensorFlow	2.37	14.35	10.25	1.27	1.84	6.83	2.08	2.79	3.07	21.49	0.68	6.44	1.43	12.06
VGG-19	PyTorch	-	-	13.88	21.71	27.63	9.65	10.74	18.54	13.81	7.68	24.98	17.40	2.93	14.06
	TensorFlow	-	-	18.20	0.92	1.16	10.58	7.34	5.06	2.74	6.92	22.88	6.37	24.12	23.56

Thank you