ATMOSPHERE
Adaptive, Trustworthy, Manageable, Orchestrated, Secure Privacy-assuring Hybrid, Ecosystem for REsilient Cloud Computing

atmosphere-eubrazil.eu

@AtmosphereEUBR

Gray-Box Models for Performance Assessment of Spark Applications

M. Lattuada, E. Gianniti, M. Hosseini, *D. Ardagna*,
A. Maros, F. Murai, A. P. Couto da Silva and J.M. Almeida

Politecnico di Milano, Italy
Universidade Federal de Minas Gerais, Brazil

Co-funded by the European Commission
Horizon 2020 - Grant #777154

Contextualizing Performance Models

- **Develop models** for:
  - Identify minimum cost configuration with a priori deadlines
  - Allow the adaptive actuation mechanisms to predict if QoS objective will be reached
  - **Assessing** (a posteriori) the main performance metrics in multi-tenancy environments
  - Evaluate if an application **run was affected by resource contention**

- **Approach**:
  - **Gray box models** based on **Machine Learning** (ML)
  - **Open source ML library**
  - Open source benchmarks

*The adoption of accurate models allows **anticipating QoS violations** and **increasing cloud services trustworthiness** by **improving** their **performance***

ML Models overview

- **Regression Models:** l1-regularized Linear Regression, Neural Network, Decision Tree, Random Forests, and Support Vector Regression

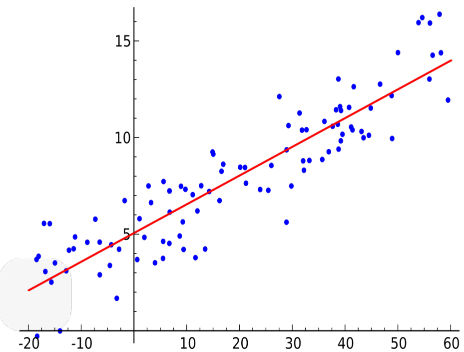

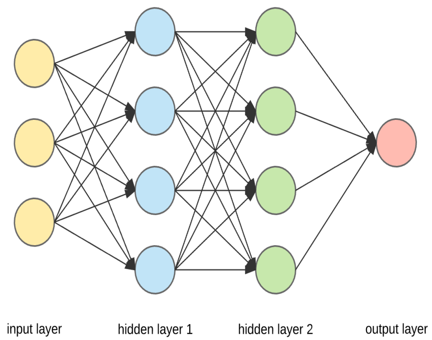

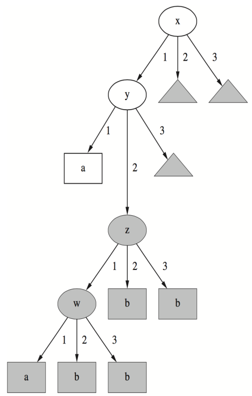

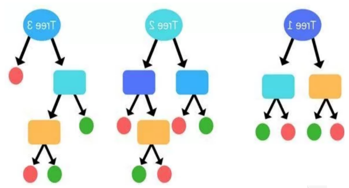

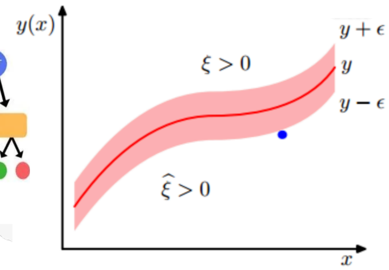

- **Hyper-parameters** optimization

ML Experiments Settings

- **Workloads:**
  - **TPC-DS** - the industry benchmark (Query 26) for data warehouse systems
  - ML benchmarks (K-means) from the **Sparkbench library**
  - **SparkDL** developed on top of Sparkbench

- **Platforms:**
  - Microsoft HDInsights on Azure - **cloud computing service**
  - IBM Power8 Cluster – **dedicated cluster**

ML Experiments Settings

- **Sampling scheme** analysis

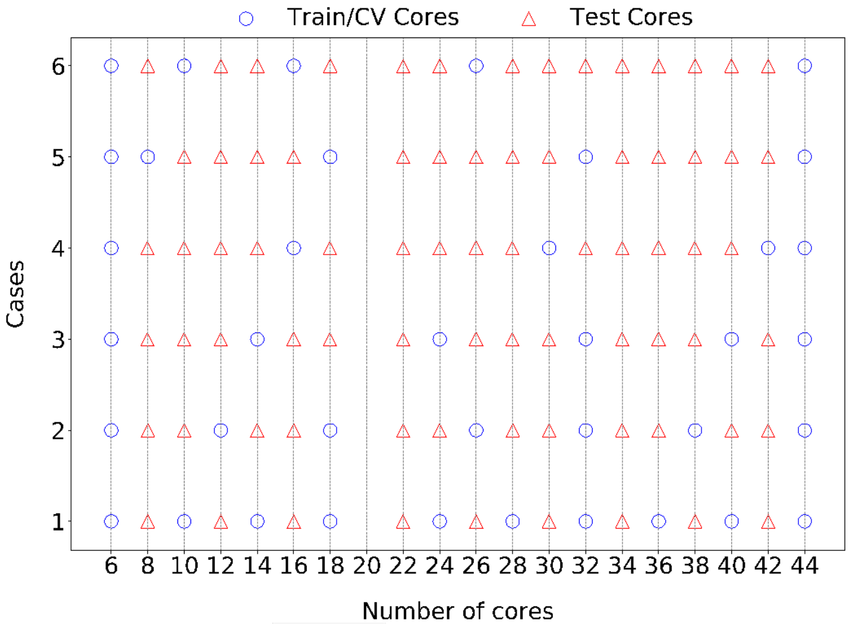

Workload	Core Interpolation		Data Extrapolation	
	Training	Test	Training	Test
Query 26 [GB]	750	750	250, 750	1000
K-means [Rows]	15	15	5, 10, 15	20
SparkDL [Images]	1500	1500	1000, 1500	2000

- **Evaluation Metric:**
  - Mean Absolute Percentage Error:

$$MAPE = \frac{100\%}{N} \sum_{k=1}^{N} \left| \frac{y_k - \hat{y}_k}{y_k} \right|$$

Comparison with State of the Art Solutions

- **Ernest Model** by Spark inventors

- Pure black-box approach:
  - Non Negative Least Square Regression (NNLS)
  - Features:
    - Ratio of data size to number of cores
    - Log of number of cores
    - Data size
    - Number of cores
    - Number of TensorFlow cores (SparkDL only)

S. Venkataraman, Z. Yang, M. Franklin, B. Recht, and I. Stoica. “Ernest: Efficient Performance Prediction for Large-Scale Advanced Analytics”. In: 13th USENIX Symposium on Networked Systems Design and Implementation. 2016, pp. 363–378.

Gray-Box Model Features

DAG structure constant across different runs

Type of Knowledge	Features
A Priori	- Ratio of data size to number of cores - Log of number of cores - Data size - Number of cores - Number of TensorFlow cores (SparkDL only)
A Posteriori	- number of tasks - max/avg time over tasks - max/avg shuffle time - max/avg number of bytes transmitted between stages - number of executor cores - inverse of number of executor cores
A Posteriori (SparkDL)	- individual TensorFlow calls execution time - inverse of total number of cores

Query 26 Results

Interpolation Scenario

	Ernest	Gray Box				
		DT	LR	NN	RF	SVR
C1	**1.50**	9.64	4.46	7.10	6.71	7.35
C2	**1.64**	9.92	8.82	5.57	12.32	20.53
C3	**1.71**	16.12	6.23	4.24	15.31	10.25
C4	**1.66**	27.05	10.62	6.09	14.37	16.92
C5	**1.59**	25.54	42.08	6.23	44.60	39.74
C6	**1.70**	11.39	35.80	6.75	41.44	68.95

Extrapolation Scenario

	Ernest	Gray Box				
		DT	LR	NN	RF	SVR
C1	**7.49**	37.13	32.91	10.23	44.16	38.08
C2	**7.44**	35.01	27.26	36.07	40.66	32.30
C3	**7.31**	32.39	36.36	20.18	47.46	38.74
C4	**7.26**	32.75	32.29	48.80	46.29	32.31
C5	**7.59**	41.93	12.41	21.90	31.50	18.11
C6	**8.02**	38.45	24.05	19.91	32.22	29.22

K-Means & SparkDL Results

Interpolation Scenario

	Ernest	Gray Box				
		DT	LR	NN	RF	SVR
C1	126.69	17.27	74.04	**11.99**	105.31	28.09
C2	148.10	21.41	83.33	**10.27**	112.77	29.13
C3	161.35	46.46	122.07	47.30	67.58	**36.14**
C4	176.52	49.66	143.33	47.60	81.70	**24.47**
C5	187.00	31.20	273.19	**29.88**	99.74	84.00
C6	159.88	**21.90**	166.57	26.66	59.40	66.75
C7	178.08	**31.59**	159.44	47.65	102.96	38.74

	Ernest	Gray Box				
		DT	LR	NN	RF	SVR
C1	10.48	5.16	5.60	**3.84**	5.87	4.12
C2	6.30	5.67	9.47	11.32	5.56	**4.66**
C3	5.71	6.40	**3.70**	5.07	8.29	4.96

Extrapolation Scenario

	Ernest	Gray Box				
		DT	LR	NN	RF	SVR
C1	167.07	26.47	20.24	59.50	26.07	**16.04**
C2	183.09	26.31	**13.15**	91.60	29.19	18.63
C3	200.83	37.28	31.86	131.54	33.47	**30.58**
C4	204.42	69.61	24.24	453.67	32.08	**18.67**
C5	221.09	26.89	26.86	**16.01**	39.72	25.29
C6	206.59	**21.19**	37.72	37.39	29.61	34.16
C7	208.37	29.44	43.14	145.81	30.21	**24.06**

	Ernest	Gray Box				
		DT	LR	NN	RF	SVR
C1	43.49	33.10	**10.73**	25.04	36.72	27.82
C2	37.39	34.94	**10.04**	17.64	35.47	32.13
C3	36.81	32.19	14.76	**9.90**	35.75	32.97

ML Models Results Summary

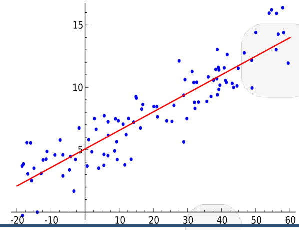

- **Gray-box** models are effective for **performance assessment** to identify **performance degradation** (about 4-25% percentage error)
  - Work better when the application **data size** is **fixed**

- Comparison with **Ernest**:in most cases, **our best models improve** Ernest considerably, especially when **few profiling configurations** are available in the **training set** and when workloads are **less regular**

Conclusions & Future work

- There is **no ML technique** which always outperforms the others, hence different techniques have to be evaluated in each scenario to choose the best model
- Study the performance of Spark applications running on **GPU- based clusters**
- Validate the models on **production environments**

Thanks for your attention…