Istituto Nazionale di Fisica Nucleare

MACHINE LEARNING @ Torino

Federica Legger, Sara Vallero

Motivation

- Provide a MLaaS platform for students and researchers
- **Yoga cluster** used by students of PhD course on Big Data and ML, and a few research projects
  - 4D fast silicon timing
  - FTS log analysis
  - MR data for earth science
- Inspired by the SWAN service@CERN, details on infrastructure given by Sara
  - No production cluster, run on best effort basis

Requirements for ML framework

- Run any ML algorithm of our choice in parallel
- Large datasets to be processed
- Multi-tenancy, so different users can request simultaneously ML pipelines
- An efficient resource management system for the cluster
- Support heterogeneous architectures CPUs, GPUs, ...

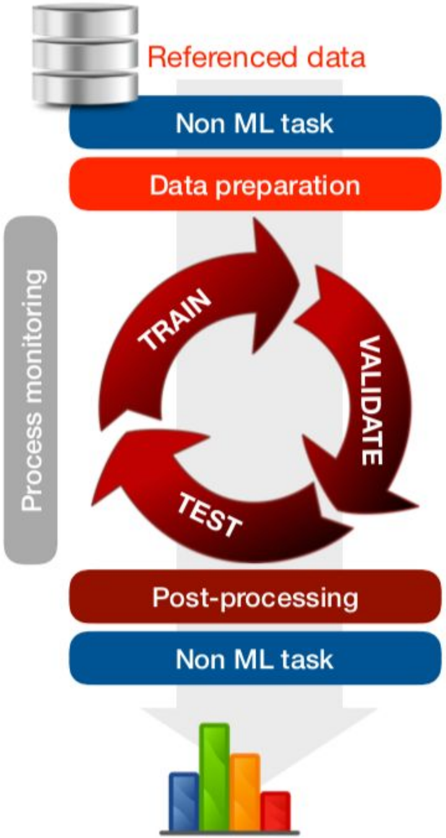

Infrastructure: requirements

Istituto Nazionale di Fisica Nucleare

- **Commodity hardware** (CPUs)
- Non-dedicated and **heterogeneous resources**:
  - Bare metal
    - 1 x 24 cores, 190 GB RAM
    - 4 x 28 cores, 260 GB RAM
  - IaaS Cloud (on premises)
    - 10 VM, 8 cores, 70 GB RAM
- Uniform application/service orchestration layer -> **Kubernetes**
- High-throughput vs. high-performance -> **Spark**
- Distributed datasets -> **HDFS**
- Elasticity: allow to scale up if there are unused resources

ML BENCHMARK FOR SCALING TESTS

- Open HEP dataset @UCI, 7GB (.csv)
- Signal (heavy Higgs) + background
- 10M MC events (balanced, 50%:50%)
  - 21 low level features
    - pt’s, angles, MET, b-tag, …
  - 7 high level features
    - Invariant masses (m(jj), m(jjj), …)

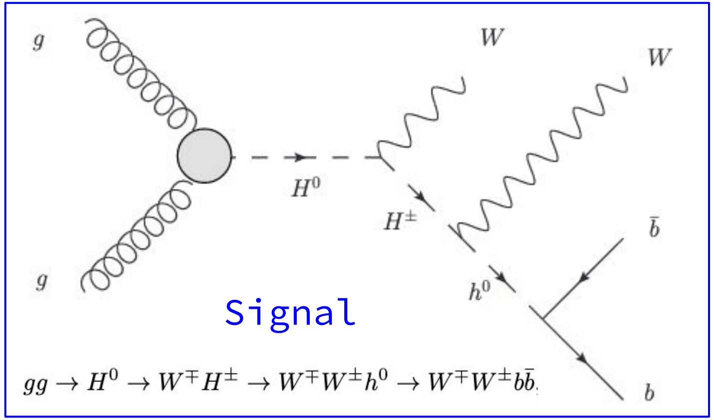

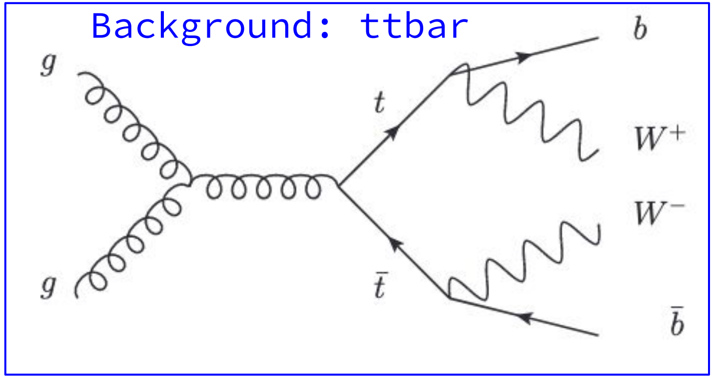

	AUC		
Technique	Low-level	High-level	Complete
BDT	0.73 (0.01)	0.78 (0.01)	0.81 (0.01)
NN	0.733 (0.007)	0.777 (0.001)	0.816 (0.004)
DN	0.880 (0.001)	0.800 (< 0.001)	0.885 (0.002)

Baldi, Sadowski, and Whiteson. “Searching for Exotic Particles in High-energy Physics with Deep Learning.” Nature Communications 5

https://archive.ics.uci.edu/ml/datasets/HIGGS

ML MODELS

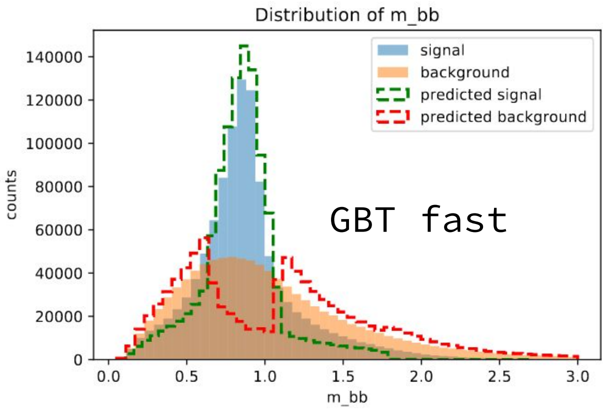
Distribution of m_bb
signal
background
predicted signal
predicted background
GBT fast
counts
m_bb
0
20000
40000
60000
80000
100000
120000
140000
0.0
0.5
1.0
1.5
2.0
2.5
3.0

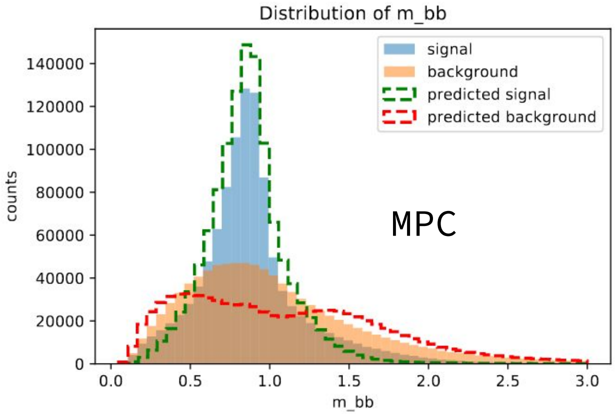
Distribution of m_bb
signal
background
predicted signal
predicted background
MPC
counts
m_bb
0
20000
40000
60000
80000
100000
120000
140000
0.0
0.5
1.0
1.5
2.0
2.5
3.0

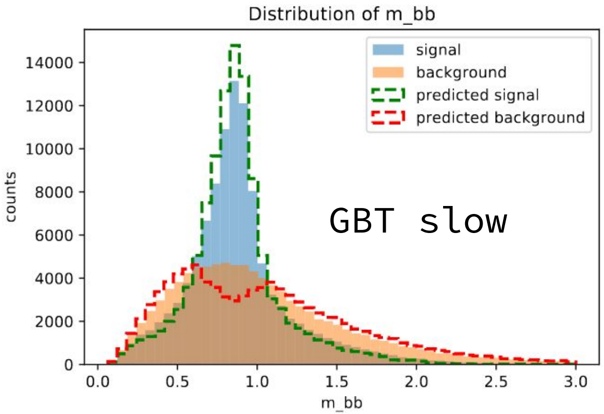
Distribution of m_bb
signal
background
predicted signal
predicted background
GBT slow
counts
m_bb
0
2000
4000
6000
8000
10000
12000
14000
0.0
0.5
1.0
1.5
2.0
2.5
3.0

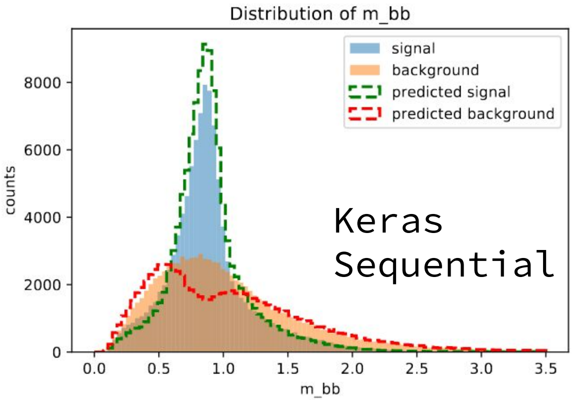
Distribution of m_bb
signal
background
predicted signal
predicted background
Keras
Sequential
counts
m_bb
0
2000
4000
6000
8000
0.0
0.5
1.0
1.5
2.0
2.5
3.0
3.5

Scaling tests #1

- Optimize #cores per executor
- Model: MLLib **MCP and GBT, 1M events**
- One machine t2-mlwn-01.to.infn.it
- In the 'literature' #cores = 5 is magic number to achieve maximum HDFS throughput

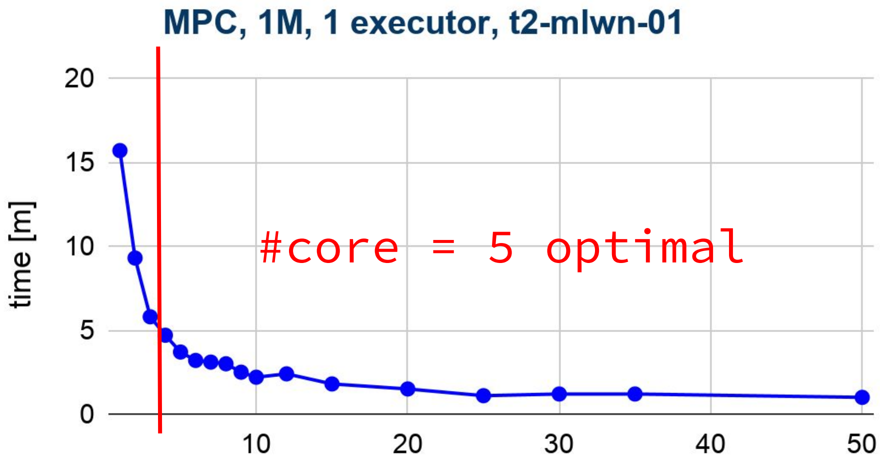

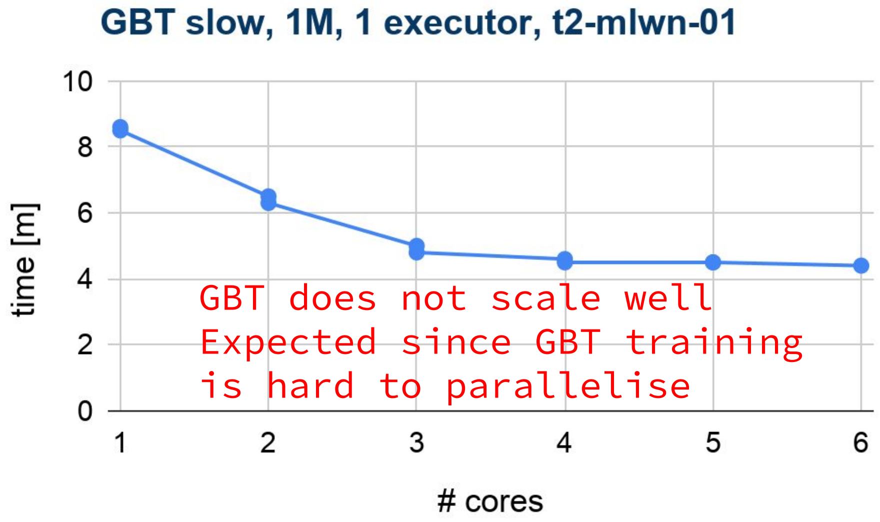

Scaling tests

Istituto Nazionale di Fisica Nucleare

- #cores per executor
- #cores per machine
- #cores in homogeneous cluster

- **Strong scaling efficiency** = time(1)/(N*time(N))
  - N = #cores

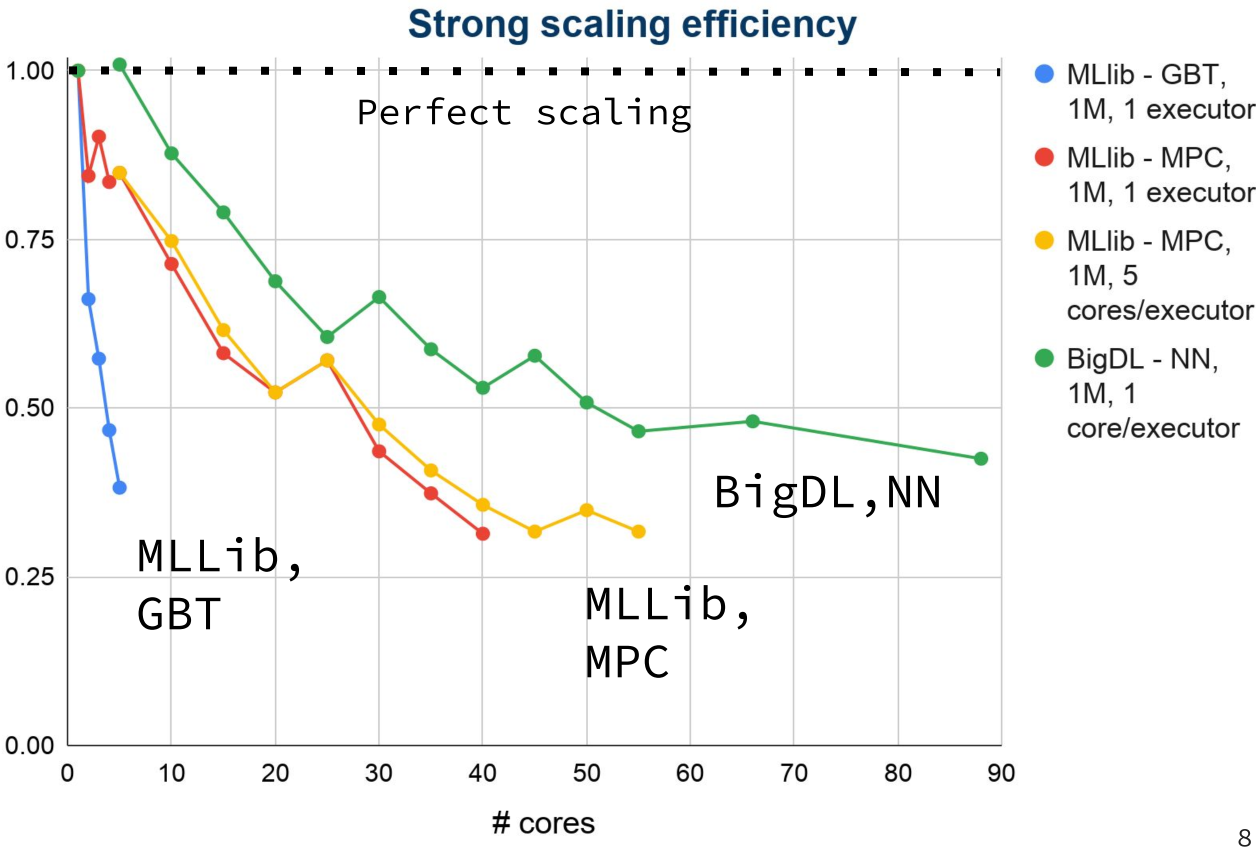

Summary

INFN
Istituto Nazionale di Fisica Nucleare

- MLaaS cluster in use in Torino
- Infrastructure runs on ‘opportunistic’ resources
- Architecture can be ‘reused’ on OCCAM

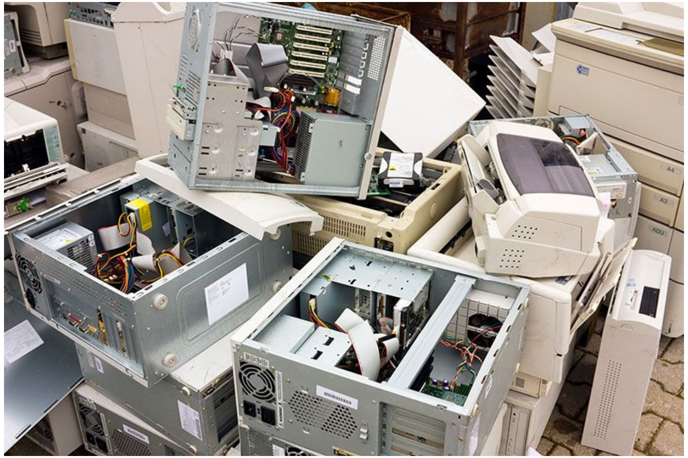

Backup

Scaling tests #2

- Optimize #executors
- Model: MLlib **MCP, 1M, 10M events**

- #cores/executor = 5
- One machine

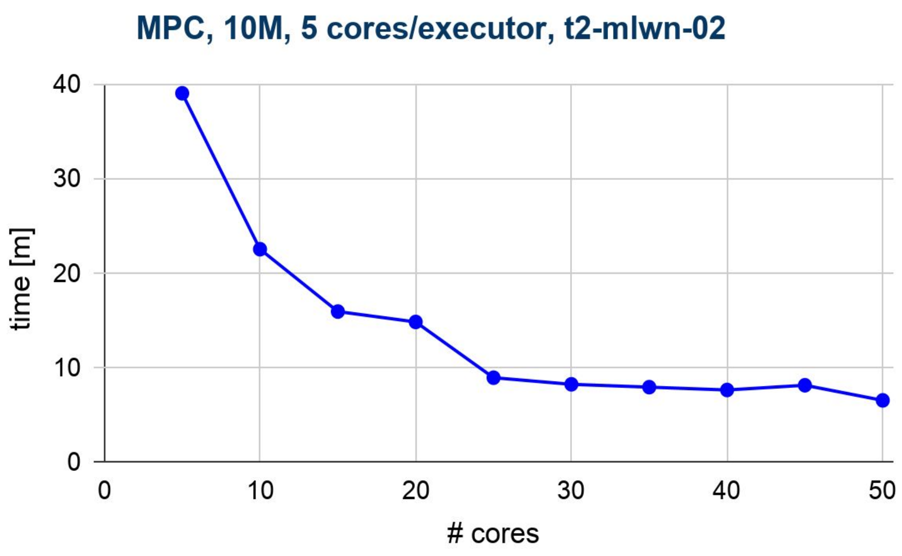

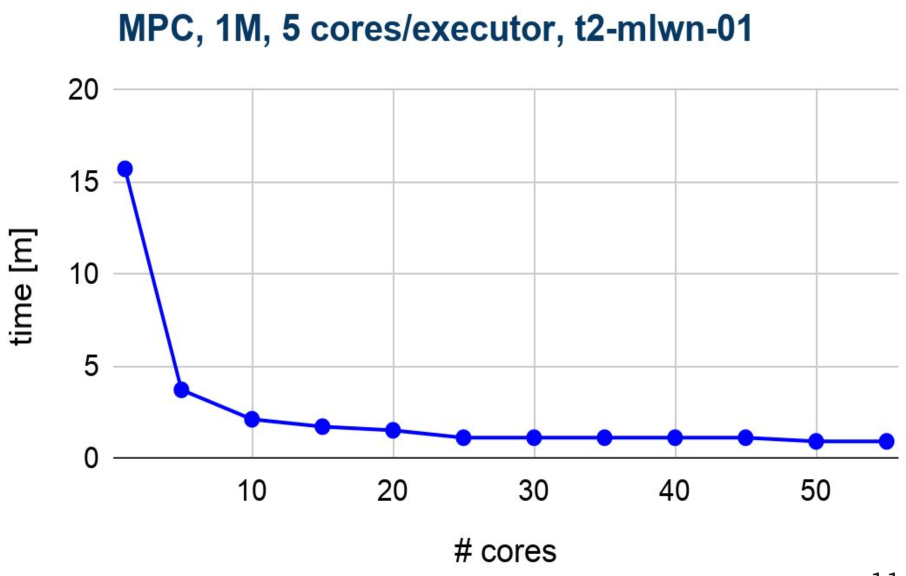

Scaling tests #3

- Scaling on homogeneous resources
  - bare metal, 4 machines with 56 cores and 260 GB

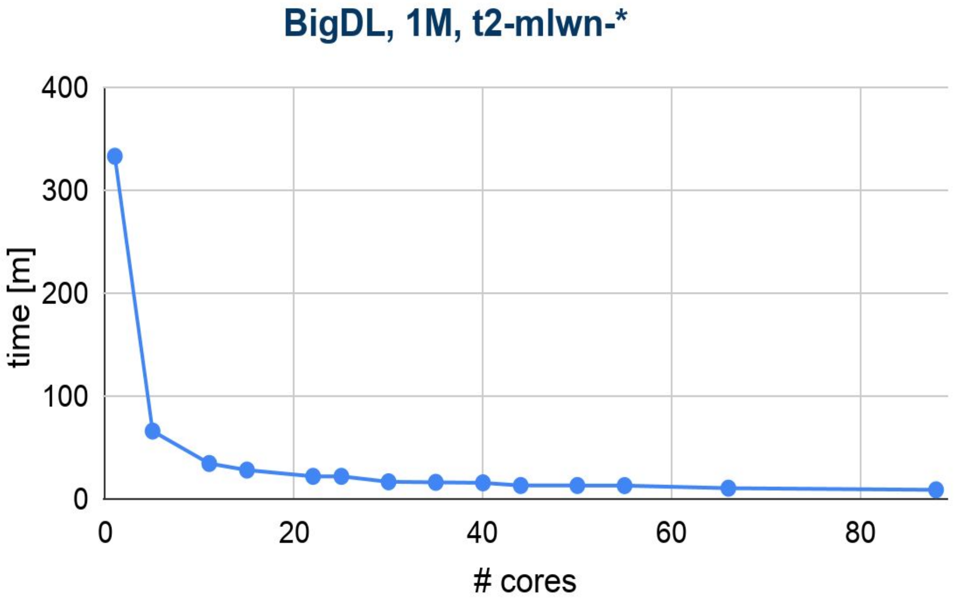

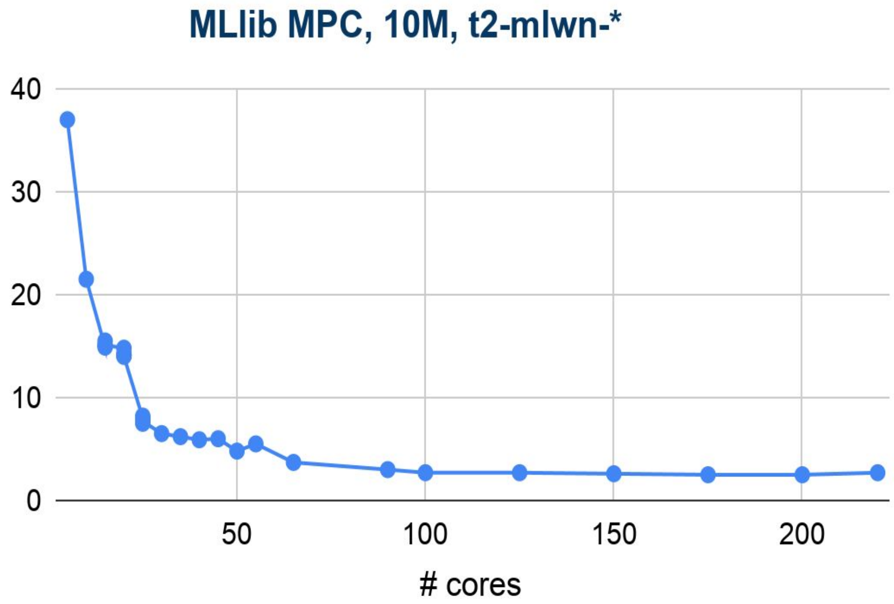

ML models && lessons learned

INFN
Istituto Nazionale di Fisica Nucleare

Model	AUC	time	# events	cores	note
MLLib GBT	82	15m	10M	25	Doesn't scale
MLLib MPC - 4 layers, 30 hidden units	74	9m	10M	25	Scales well, can't build complex models
Keras Sequential - 1 layer, 100 hidden units	81	18m	1M	25	No distributed training, cannot process 10M events
BigDL Sequential - 2 layers, 300 hidden units	86	3h15m	10M	88	1 core/executor required