



dataset repositories

Machine Learning for the Big Data Benchmark Repository and Performance Analysis

ALOJA Project

The ALOJA research project is an initiative from the Barcelona Supercomputing Center (BSC) to explore new hardware architectures for Big Data processing. One of the main goals of the project is to produce a systematic study of SW and HW configuration and deployment options; where we are analyzing the cost-effectiveness of the different cloud services as well as on-premise hardware, both commodity and up-scale.

ALOJA + Machine Learning

ALOJA-ML is the set of Machine Learning autonomous scripts, prepared to run on the ALOJA project. Also, ALOJA-ML takes care of performing data mining, models and prediction on the datasets generated in the ALOJA project.

Here you can find the links to the GitHub pages for ALOJA and ALOJA-ML, and the Barcelona Supercomputing Center. Further, you can find the project publications, also the structured datasets for those publications.

ALOJA & ALOJA-ML Publications

-  [ALOJA Website](#)
-  [ALOJA-ML GitHub Page](#)
-  [Barcelona Supercomputing Center](#)

David Buchaca, Josep Ll. Berral, David Carrera.

Automatic Generation of Workload Profiles using Unsupervised Learning Pipelines. IEEE Transactions on Networks and Systems Management (TNSM), vol.15 issue.1 pp.142-155 (2017). ISSN 1932-4537. Open Access.

Josep Ll. Berral, Nicolas Poggi, David Carrera, Aaron Call, Rob Reinauer, Daron Green. **ALOJA: A Framework for Benchmarking and Predictive Analytics in Big Data Deployments.** IEEE Transactions on Emerging Topics in Computing (TETC), vol.5 issue.4 pp.480-493 (2017). ISSN 2168-6750. arXiv:1511.02037.

Nicolas Poggi, Josep Ll. Berral, David Carrera. **ALOJA: a Benchmarking and Predictive Platform for Big Data Performance Analysis.** The Sixth Workshop on Big Data Benchmarking (6th WBDB). June 16-17, 2015 in Toronto, Canada.

Nicolas Poggi, Josep Ll. Berral, David Carrera, Aaron Call, Rob Reinauer, Nikola Vujic, Daron Green, José Blakeley, Fabrizio Gagliardi. **From Performance Profiling to Predictive Analytics while Evaluating Hadoop Cost-Efficiency in ALOJA.** The IEEE International Conference on Big Data (IEEE BigData 2015), Santa Clara (CA), USA, Oct. 29-Nov. 1 2015.

Josep Ll. Berral, Nicolas Poggi, David Carrera, Aaron Call, Rob Reinauer, Daron Green. **ALOJA-ML: A Framework for Automating Characterization and Knowledge Discovery in Hadoop Deployments.** The 21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2015), Sydney, Australia, August 10-13 2015. arXiv:1511.02030.

Nicolas Poggi, David Carrera, Aaron Call, Rob Reinauer, Nikola Vujic, Daron Green and Jose Blakeley, et al. **ALOJA: a Systematic Study of Hadoop Deployment Variables to Enable Automated Characterization of Cost-Effectiveness.** IEEE BigData 2014. 27-30 Oct. 2014. Washington DC, USA.

far as you cite us (check each dataset references for the corresponding publication). If you obtain good (or interesting) results by applying your methods and techniques when predicting elements on our datasets, you can communicate with us to publish the results in our method rankings, providing the details, notes, evidences or link to your corresponding publication.

Here you can find the [LEGAL NOTICE](#) for BSC-CNS and content in this site.

ALOJA Spark Time-Series Dataset

This dataset comprises 900 executions of 30 different Spark applications from TPCx-BB (BigBench) benchmark, with different types of workloads (NLP, SQL, MapReduce, Machine Learning, UDTFs...), different data types (Structured, Semi-Structured and Un-structured data), and different data scales (1,10 and 100GB). All the jobs were run in the Microsoft's Azure cloud using Spark 2 as the engine, HDInsight PaaS to spawn the spark clusters, running a 16-slave node cluster, and data was stored in the Azure Data Lake Store of Azure. Reference: the **Workload Profiles** paper at **IEEE-TNSM'17**.

Dataset details:	
Number of entries	900 executions, 30 different Spark applications (7121338 time entries)
Notes	The 'simplified' files contain the aggregate information for all nodes. The 'complete' files contain the information separated for each headnode and datanode. Some attributes may have missing values, marked as -1 when Not Available.
Dataset features:	
Time Attributes	timestamp, interval, instant
Execution Attributes	job_name, disk (type), query_name, cached, platform, engine, query (number)
Performance Attributes	X.{usr, nice, sys, iowait, steal, irq, soft, guest, gnice, idle} (cpu percent), kbmemb{used, free}, X.{memused, commit} (mem percent), kb{buffers, cached, commit, active, inactive, dirty, anonpg, slab, kstack, pgtbl, vmused}, {rx, tx}pck.s, {rx, tx}Kb.s, {rx, tx}cmp.s, {rx, tx}cst.s, X.ifutil (iface percent), tps, {rd, wr}_sec.s, avgrq.sz, avgqu.sz, await, svctm, X.util (disk percent)
Download Dataset »	




ALOJA Hadoop Time-Series Dataset



This dataset contains 182 series from Hadoop executions from the Intel Hi-Bench benchmark suite, with Map-Reduce algorithms for sorting, word-counting, machine learning, input-output stress testing, etc. All the jobs have been running in on-premise infrastructures, with similar Hadoop configurations. Reference: the **ALOJA** paper at **IEEE-TETC'17**

Dataset details:	
Number of entries	182 executions (368683 time entries)
Notes	Contains Hadoop execution logs per time unit, indicating timestamp, number of workers, and consumed resources. Some of these attributes may have missing values, marked as -1 when Not Available.
Dataset features:	
Time Attributes	instant, date
Execution	id_JOB_job_status, id_exec, job_name, JOBD, bench

<http://aloja.bsc.es/>

AUG NOV DEC
◀ 19 ▶
2017 2019 2020

[14 captures](#)
 9 Oct 2015 - 19 Nov 2019

Performance Attributes: pc.{user, system, iowait}, kbmemused, {rx, tx}pck.s, tps, rtps, wtps

[Download Dataset >](#)

ALOJA Hadoop Dataset v6

This dataset contains traces of Hadoop executions. Slice of Aloja Dataset v5, including the aggregated resource performance per execution. Executions include, at least, performance information for CPU and Memory. Some executions miss Network or Disk information. Reference: **ALOJA-ML** paper at **KDD'15**.

Dataset details:	
Number of entries	33147 executions
Notes	Some of these attributes (valid, filter, outlier) are not completely reliable, and are based on automatic filtering of executions. Beware when using them. Comp (compression) is coded as [0: None, 1: ZLIB, 2: BZIP2, 3: Sappy]
Dataset features:	
Execution Attributes	ID, Start.Time, End.Time, Valid, Filter, Outlier, Perf.Details, Run.Num.
Configuration Attributes	Benchmark, Net, Disk, Bench.Type, Maps, IO.SFac, Rep, IO.FBuf, Comp, Blk.Size, Hadoop.Version, Exec.Type, Datasize, Scale.Factor, Java.XMS, Java.XMX
Cluster Attributes	Cluster (ID), Cl.Name, Service.Type, Datanodes, Headnodes, VM.Size, VM.OS, VM.Cores, VM.RAM, Provider
Cost values	Cost.Remote, Cost.SSD, Cost.IB, Cost.Hour
Time Performance Attributes	Exe.Time
Resource Performance Attributes	<p>CPU features (percentage from single CPUs): {avg, max, min, stdev_pop, var_pop}.user, {avg, max, min, stdev_pop, var_pop}.nice, {avg, max, min, stdev_pop, var_pop}.system, {avg, max, min, stdev_pop, var_pop}.iowait, {avg, max, min, stdev_pop, var_pop}.steal, {avg, max, min, stdev_pop, var_pop}.idle</p> <p>Memory features: {avg, max, min, stdev_pop, var_pop}.kbmemfree, {avg, max, min, stdev_pop, var_pop}.kbmemused, {avg, max, min, stdev_pop, var_pop}.memused (percentage from total mem), {avg, max, min, stdev_pop, var_pop}.kbbuffers, {avg, max, min, stdev_pop, var_pop}.kbcached, {avg, max, min, stdev_pop, var_pop}.kbcommit, {avg, max, min, stdev_pop, var_pop}.commit, {avg, max, min, stdev_pop, var_pop}.kbactive, {avg, max, min, stdev_pop, var_pop}.kbinact</p> <p>Network features: {avg, max, min, stdev_pop, var_pop, sum}.rxpck.s, {avg, max, min, stdev_pop, var_pop, sum}.txpck.s, {avg, max, min, stdev_pop, var_pop, sum}.rxkB.s, {avg, max, min, stdev_pop, var_pop, sum}.txkB.s, {avg, max, min, stdev_pop, var_pop, sum}.rxcmp.s, {avg, max, min, stdev_pop, var_pop, sum}.txcmp.s, {avg, max, min, stdev_pop, var_pop, sum}.rxmct.s</p> <p>Disk features: {avg, max, min, stdev_pop, var_pop, sum}.rd_sec.s, {avg, max, min, stdev_pop, var_pop, sum}.wr_sec.s, {avg, max, min, stdev_pop, var_pop, sum}.rq_sz, {avg, max, min, stdev_pop, var_pop, sum}.qu_sz, {avg, max, min, stdev_pop, var_pop, sum}.await, {avg, max, min, stdev_pop, var_pop, sum}.util, {avg, max, min, stdev_pop, var_pop, sum}.svctm</p>
Download Dataset >	

Reference: **ALOJA-ML** paper at **KDD'15**.

Dataset details:	
Number of entries	43649 executions
Notes	Some of these attributes (valid, filter, outlier) are not completely reliable, and are based on automatic filtering of executions. Beware when using them. Comp (compression) is coded as [0: None, 1: ZLIB, 2: BZIP2, 3: Sappy]
Dataset features:	
Execution Attributes	ID, Valid, Filter, Outlier
Configuration Attributes	Benchmark, Net, Disk, Bench.Type, Maps, IO.SFac, Rep, IO.FBbuf, Comp, Blk.Size, Hadoop.Version
Cluster Attributes	Cluster (ID), Cl.Name, (Service) Type, Datanodes, Headnodes, VM.Size, VM.OS, VM.Cores, VM.RAM, Provider
Time Performance Attributes	Exe.Time
Download Dataset >	

ALOJA Hadoop Dataset RAW

This dataset contains raw traces of Hadoop executions, from configurations to SAR records. Reference: **ALOJA** paper at **IEEE-TETC'17**.

Dataset details:	
Number of entries	>50K executions, >800M of records in profiling time series.
Notes	Files contain records of execution results and execution raw traces taken with SAR, VMSTAT and other profiling tools. The full data-set occupies more than 2 TB.
ALOJA Files:	
HDI_JOB_details	hdi_job_details_id, id_exec, job_id, bytes_read, bytes_written, committed_heap_bytes, cpu_milliseconds, failed_maps, failed_reduces, failed_shuffle, file_bytes_read, file_bytes_written, file_large_read_ops, file_read_ops, file_write_ops, finished_maps, finish_time, gc_time_millis, job_priority, launch_time, map_input_records, map_output_records, mb_millis_maps, merged_map_outputs, millis_maps, other_local_maps, physical_memory_bytes, slots_millis_maps, spilled_records, split_raw_bytes, submit_time, total_launched_maps, total_maps, total_reduces, user, vcores_millis_maps, virtual_memory_bytes, wasb_bytes_read, wasb_bytes_written, wasb_large_read_ops, wasb_read_ops, wasb_write_ops, job_name, records_written, bad_id, combine_input_records, combine_output_records, connection, io_error, map_output_bytes, map_output_materialized_bytes, mb_millis_reduces, millis_reduces, rack_local_maps, reduce_input_groups, reduce_input_records, reduce_output_records, reduce_shuffle_bytes, wrong_length, wrong_map, wrong_reduce, total_launched_reduces, shuffled_maps, slots_millis_reduces, vcores_millis_reduces, checksum, num_failed_maps, hdfs_bytes_read,

14 captures

9 Oct 2015 - 19 Nov 2019

	finish_time, job_priority, user, total_maps, failed_maps, finished_maps, total_reduces, failed_reduces, launched map tasks, rack-local map tasks, launched reduce tasks, slots_millis_maps, slots_millis_reduces, data-local map tasks, file_bytes_written, file_bytes_read, hdfs_bytes_written, hdfs_bytes_read, bytes read, bytes written, spilled records, split_raw_bytes, map input records, map output records, map input bytes, map output bytes, map output materialized bytes, reduce input groups, reduce input records, reduce output records, reduce shuffle bytes, combine input records, combine output records
JOB_dbscan	id, bench, job_offset, metric_x, metric_y, TASK_TYPE, id_exec, centroid_x, centroid_y
clusters	id_cluster, name, cost_hour, type, link, datanodes, headnodes, vm_size, vm_OS, vm_cores, vm_RAM, description, provider, cost_remote, cost_SSD, cost_IB
execs	id_exec, id_cluster, exec, bench, exe_time, start_time, end_time, net, disk, bench_type, maps, iosf, replication, iofilebuf, comp, blk_size, zabbix_link, hadoop_version, valid, filter, outlier, perf_details, exec_type, datasize, scale_factor, JAVA_XMS, JAVA_XMX, run_num
hosts	id_host, host_name, id_cluster, role, cost_remote, cost_SSD, cost_IB
precal_cpu_metrics	id_exec, avg%user, max%user, min%user, stddev_pop%user, var_pop%user, avg%nice, max%nice, min%nice, stddev_pop%nice, var_pop%nice, avg%system, max%system, min%system, stddev_pop%system, var_pop%system, avg%iowait, max%iowait, min%iowait, stddev_pop%iowait, var_pop%iowait, avg%steal, max%steal, min%steal, stddev_pop%steal, var_pop%steal, avg%idle, max%idle, min%idle, stddev_pop%idle, var_pop%idle
precal_disk_metrics	id_exec, DEV, avgtps, maxtps, mintps, avgrd_sec/s, maxrd_sec/s, minrd_sec/s, stddev_poprd_sec/s, var_poprd_sec/s, sumrd_sec/s, avgwr_sec/s, maxwr_sec/s, minwr_sec/s, stddev_popwr_sec/s, var_popwr_sec/s, sumwr_sec/s, avgrq_sz, maxrq_sz, minrq_sz, stddev_poprq_sz, var_poprq_sz, avgqu_sz, maxqu_sz, minqu_sz, stddev_popqu_sz, var_popqu_sz, avgawait, maxawait, minawait, stddev_popawait, var_popawait, avg%util, max%util, min%util, stddev_pop%util, var_pop%util, avgsvctm, maxsvctm, minsvctm, stddev_popsvctm, var_popsvctm
precal_memory_metrics	id_exec, DEV, avgkbmemfree, maxkbmemfree, minkbmemfree, stddev_popkbmemfree, var_popkbmemfree, avgkbmemused, maxkbmemused, minkbmemused, stddev_popkbmemused, var_popkbmemused, avg%memused, max%memused, min%memused, stddev_pop%memused, var_pop%memused, avgkbbuffers, maxkbbuffers, minkbbuffers, stddev_popkbbuffers, var_popkbbuffers, avgkbcached, maxkbcached, minkbcached, stddev_popkbcached, var_popkbcached, avgkbcommit, maxkbcommit, minkbcommit, stddev_popkbcommit, var_popkbcommit, avg%commit, max%commit, min%commit, stddev_pop%commit, var_pop%commit, avgkbactive, maxkbactive, minkbactive, stddev_popkbactive, var_popkbactive, avgkbinact, maxkbinact, minkbinact, stddev_popkbinact, var_popkbinact
precal_network_metrics	id_exec, IFACE, avgrxpck/s, maxrxpck/s, minrxpck/s, stddev_poprxpck/s, var_poprxpck/s, sumrxpck/s, avgtxpck/s, maxtxpck/s, mintxpck/s, stddev_poptxpck/s, var_poptxpck/s, sumtxpck/s, avgrxB/s, maxrxkB/s, minrxkB/s, stddev_poprxkB/s, var_poprxkB/s, sumrxkB/s, avgtxkB/s, maxtxkB/s, mintxkB/s, stddev_poptxkB/s, var_poptxkB/s, sumtxkB/s, avgrxcmp/s, maxrxcmp/s, minrxcmp/s, stddev_poprxcmp/s, var_poprxcmp/s,

[Download Dataset »](#)

ALOJA_logs Files:

BWM	id_BWM, id_exec, host, unix_timestamp, iface_name, bytes_out, bytes_in, bytes_total, packets_out, packets_in, packets_total, errors_out, errors_in
BWM2	id_BWM, id_exec, host, unix_timestamp, iface_name, bytes_out/s, bytes_in/s, bytes_total/s, bytes_in, bytes_out, packets_out/s, packets_in/s, packets_total/s, packets_in, packets_out, errors_out/s, errors_in/s, errors_in, errors_out

[Download Files »](#)

HDI_JOB_tasks	hdi_job_task_id, job_id, task_id, bytes_read, bytes_written, committed_heap_bytes, cpu_milliseconds, failed_shuffle, file_bytes_read, file_bytes_written, file_read_ops, file_write_ops, gc_time_millis, map_input_records, map_output_records, merged_map_outputs, physical_memory_bytes, spilled_records, split_raw_bytes, task_error, task_finish_time, task_start_time, task_status, task_type, virtual_memory_bytes, wasb_bytes_read, wasb_bytes_written, wasb_large_read_ops, wasb_read_ops, wasb_write_ops, file_large_read_ops, records_written, map_output_bytes, map_output_materialized_bytes, combine_input_records, combine_output_records, id_exec, reduce_input_groups, reduce_output_groups, reduce_shuffle_bytes, reduce_input_records, reduce_output_records, shuffled_maps, bad_id, io_error, wrong_length, connection, wrong_map, wrong_reduce, checksum, num_failed_maps, hdfs_bytes_read, hdfs_bytes_written, hdfs_large_read_ops, hdfs_large_write_ops, hdfs_read_ops, hdfs_write_ops, job_name, created_files, deserialize_errors, failed_reduces, finished_maps, job_priority, launch_time, mb_millis_maps, mb_millis_reduces, millis_maps, millis_reduces, num_killed_maps, num_killed_reduces, other_local_maps, rack_local_maps, records_in, records_out_intermediate, skewjoinfollowupjobs, slots_millis_maps, slots_millis_reduces, submit_time, total_launched_maps, total_launched_reduces, total_maps, total_reduces, user, vcores_millis_maps, vcores_millis_reduces, data_local_maps
JOB_status	id_job_job_status, id_exec, job_name, jobid, date, maps, shuffle, merge, reduce, waste
JOB_tasks	id_job_job_tasks, id_exec, job_name, jobid, taskid, task_type, task_status, start_time, finish_time, shuffle_time, sort_time, bytes read, bytes written, file_bytes_written, file_bytes_read, hdfs_bytes_written, hdfs_bytes_read, spilled records, split_raw_bytes, map input records, map output records, map input bytes, map output bytes, map output materialized bytes, reduce input groups, reduce input records, reduce output records, reduce shuffle bytes, combine input records, combine output records

[Download Files »](#)

SAR_block_devices	id_SAR_block_devices, id_exec, host, interval, date, DEV, tps, rd_sec/s, wr_sec/s, avgrq-sz, avgqu-sz, await, svctm, %util
SAR_cpu	id_SAR_cpu, id_exec, host, interval, date, CPU, %user, %nice, %system, %iowait, %steal, %idle
SAR_interrupts	id_SAR_interrupts, id_exec, host, interval, date, INTR, intr/s"
SAR_io_paging	id_SAR_io_paging, id_exec, host, interval, date, pgpgin/s, pgpgout/s, fault/s, majflt/s, pgfree/s, pgscank/s, pgscand/s, pgsteal/s, %vmeff

SAR_memory	id_SAR_memory, id_exec, host, interval, date, frmpg/s, bufpg/s, campg/s
SAR_memory_util	id_SAR_memory_util, id_exec, host, interval, date, kbmemfree, kbmemused, %memused, kbbuffers, kbcached, kbcommit, %commit, kbactive, kbinact, kbdirty
SAR_net_devices	id_SAR_net_devices, id_exec, host, interval, date, IFACE, rxpck/s, txpck/s, rxkB/s, txkB/s, rxcmp/s, txcmp/s, rxmcast/s, %ifutil
SAR_net_errors	id_SAR_net_errors, id_exec, host, interval, date, IFACE, rxerr/s, txerr/s, coll/s, rxdrop/s, txdrop/s, txcarrr/s, rxfram/s, rxfifo/s, txfifo/s
SAR_net_sockets	id_SAR_net_sockets, id_exec, host, interval, date, totsck, tcpsck, udpsck, rawsck, ip-frag, tcp-tw
SAR_swap	id_SAR_swap, id_exec, host, interval, date, kbswpfree, kbswpused, %swpused, kbswpcad, %swpcad
SAR_swap_util	id_SAR_swap_util, id_exec, host, interval, date, pswpin/s, pswpout/s
SAR_switches	id_SAR_switches, id_exec, host, interval, date, proc/s, cswch/s
VMSTATS	id_VMSTATS, id_exec, host, time, r, b, swpd, free, buff, cache, si, so, bi, bo, in, cs, us, sy, id, wa, st
Download Files »	

Acknowledgements



Contact

- C\ Jordi Girona 31, 08034 Barcelona, Spain
- Telf: (+34) 93 413 40 78
Fax: (+34) 93 413 77 21
- Mail: josep.berral [at] bsc.es

Barcelona Supercomputing Center

Barcelona Supercomputing Center-Centro Nacional de Supercomputación (BSC-CNS) is the national supercomputing centre in Spain. We specialise in high performance computing (HPC) and manage MareNostrum, one of the most powerful supercomputers in Europe, located in the Torre Girona chapel.

[Barcelona Supercomputing Center - ALOJA-ML](#) © 2015 - 2019 - All Rights Reserved

Template by [OS Templates](#)