

FTS Log Analysis in Atlas - Status and Plans



Mayank Sharma
(on behalf of the Operational Intelligence collaborators)

NLP components in OpInt Framework

- Abstract classes and methods
- Provide an interface between the framework and various ML approaches for log analysis
- ML experts arrange their code into 3 classes provided by the framework
 - Tokenization
 - Vectorization
 - Clusterization
- ML Experts are free to use any libraries
- OpInt Framework services can train/load models, evaluate results on demand (UI) or
- [Specification document](#)
- [OpInt F/w code](#)

NLP components in Oplnt Framework

- Support for Jupyter Notebooks
- ML experts can use their NLP components of the Framework in jupyter notebooks
- Good to test implementation of ML approach within the Framework

```
10 def train_kmeans(dataset, k, ft_col='features', distance="cosine", initSteps=100,
11                 tol=0.00001, maxIter=50, save_path=None, mode="new", log_path=None):
12     """Train K-Means model.
13
14     -- params:
15     dataset (pyspark.sql.dataframe.DataFrame): data frame with a vector column with features for the kmeans algorithm
16     k (int): number of clusters
17     ft_col (string): name of the features column
18     distance ("euclidean" or "cosine"): distance measure for the kmeans algorithm
19     tol (int): tolerance for kmeans algorithm convergence
20     maxIter (int): maximum number of iterations for the kmeans algorithm
21     save_path (string): where to save trained kmeans model
22     mode ("new" or "overwrite"): whether to save new file or overwrite pre-existing one.
23     log_path (string): where to save optimization stats. Default None (no saving)
24
25     Returns:
26     model_fit (pyspark.ml.clustering.KMeansModel): trained K-Means model
27     """
```

FTS Log Analysis: Data

- FTS data fetched from HDFS

t_error_message	src_hostname	dst_hostname	tr_timestamp_complete
DESTINATION OVERWRITE Result (Neon): Could not read status line: Secure connection truncated after 1 attempts	ccdavatlas.in2p3.fr	dynafed-atlas.cern.ch	2020-10-04 17:42:36.041
TRANSFER ERROR: Copy failed with mode 3rd pull, with error: copy (Neon): Could not parse redirect destination URL	sdrm.t1.grid.kiae.ru	eosatlas.cern.ch	2020-10-04 17:41:54.533
DESTINATION OVERWRITE Result (Neon): Could not read status line: Secure connection truncated after 1 attempts	t2-dpm-01.na.infn.it	dynafed-atlas.cern.ch	2020-10-04 17:43:33.281
[gfalt_copy_file][perform_copy][gfal_http_copy] DESTINATION OVERWRITE [gfal_http_copy_overwrite] [gfal_http_exists] [gfal_http_access][davix2gliberr] Failure (Neon): Server certificate verification failed: certificate has expired after 1 attempts	se03.esc.qmul.ac.uk	dynafed.stfc.ac.uk	2020-10-04 17:44:48.205
DESTINATION OVERWRITE Result (Neon): Could not read status line: Secure connection truncated after 1 attempts	atlas.dcache.nikhef.nl	dynafed-atlas.cern.ch	2020-10-04 17:42:59.003
DESTINATION OVERWRITE Result (Neon): Could not read status line: Secure connection truncated after 1 attempts	sdrm.t1.grid.kiae.ru	dynafed-atlas.cern.ch	2020-10-04 17:41:58.083

FTS Log Analysis: Results

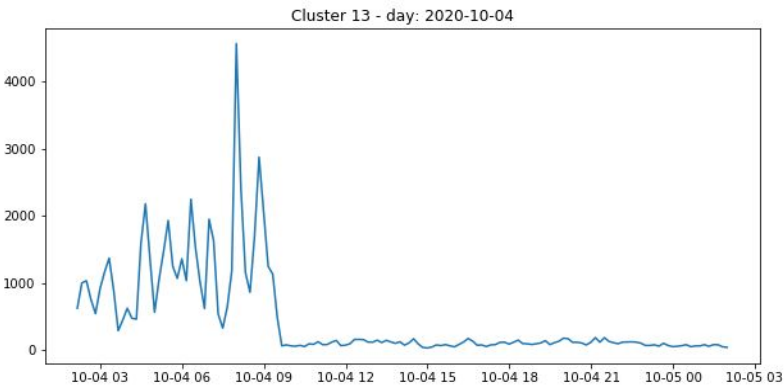



- Cluster Summary
- Word Clouds
- Time Plots

FTS Log Analysis: Cluster Summary

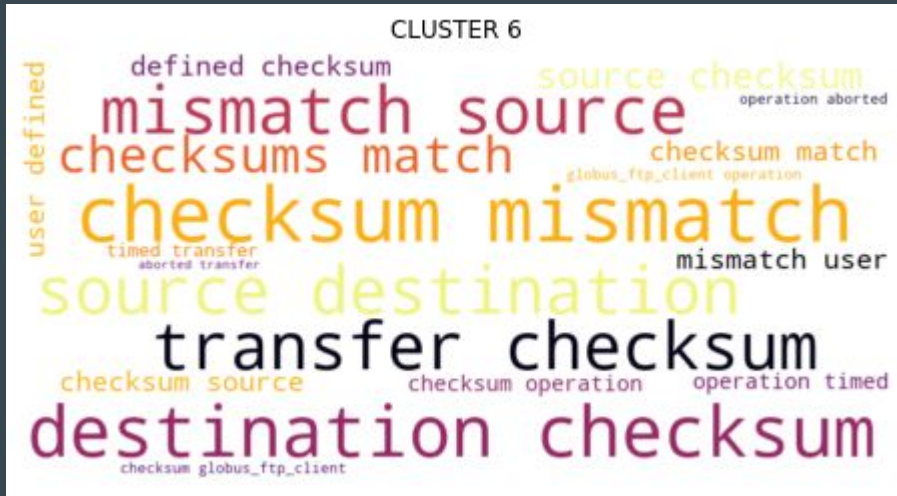
Full results at: [sample_app_results](#)

Progressive ID	ID	n_messages	unique_strings	unique_patterns	top_5_msg	top_5_src	top_5_dst
4	13	65574	6760	73	msg: 'transfer globus_ftp_client the server responded with an error 500 command failed internal server error', 'n': 45867, 'n_perc': 0.6995	src: 'NDGF-T1_SE', 'n': 5029, 'n_perc': 0.0767	dst: 'GRIF_LAL_SE_LHCB', 'n': 45690, 'n_perc': 0.6968
					msg: 'transfer globus_ftp_client the server responded with an error 500 command failed open/create [error] server responded with an error [3021] unable to get quota space quota not defined or exhausted \\\$FILE_PATH disk quota exceeded', 'n': 8656, 'n_perc': 0.132	src: 'SARA-MATRIX_SE', 'n': 4658, 'n_perc': 0.071	dst: 'USCMS-FNAL-WC1_SE_CMS', 'n': 6558, 'n_perc': 0.1
					msg: '[gfalt_copy_file][perform_copy][gridftp_plugin_filecopy][gridftp_plugin_filecopy] transfer globus_ftp_client the server responded with an error 500 command failed open/create [error] server responded with an error [3021] unable to get quota space quota not defined or exhausted \\\$FILE_PATH disk quota exceeded', 'n': 2403, 'n_perc': 0.0366	src: 'NIKHEF-ELPROD_SE_167_ATLAS', 'n': 4564, 'n_perc': 0.0696	dst: 'CERN-PROD_SE_CMS', 'n': 4680, 'n_perc': 0.0714
					msg: 'transfer globus_ftp_client the server responded with an error 500 command failed ipc failed while attempting to perform request', 'n': 2175, 'n_perc': 0.0332	src: 'RAL-LCG2_SE_1_ATLAS', 'n': 4432, 'n_perc': 0.0676	dst: 'IFCA-LCG2_SE_CMS', 'n': 1398, 'n_perc': 0.0213
					msg: 'source checksum globus_ftp_client the server responded with an error 500 error commands denied', 'n': 1489, 'n_perc': 0.0227	src: 'RRC-KI-T1_SE_1_ATLAS', 'n': 2845, 'n_perc': 0.0434	dst: 'UKI-SCOTGRID-GLASGOW_SE_0_ATLAS', 'n': 1132, 'n_perc': 0.0173

FTS Log Analysis: Time Plots



FTS Log Analysis: Word Cloud



FTS Log Analysis: Practical Implications

- Luca recently presented Atlas Software and Computing Week
 - [Slides](#)
- Detected clusters dominated by certain sites based on the output of his ML approach
- Found a corresponding GGUS ticket for the site

FTS Log Analysis: Approach

Start from error messages and try to group them based on meaning. The extra information is considered in a second phase.

In order to do that, we can distinguish two steps:

- **vectorization** (*tokenization + language model^(*)*): first transform text to numeric
- **clustering^(#)**: then try to group numeric representations of messages

Current strategy: online re-training

Train a vectorization model once on a big dataset (possibly updating once in a while).

Re-train a clustering algorithm online every time new data come in (perhaps 1 day time windows)

Implementation in OpInt Framework

- Luca provided implementation of his algorithms corresponding to the abstract NLP classes of the OpInt Framework
- PySpark based approach
 - Tokenization
 - Vectorization
 - Clustering
- Jupyter Notebook
- Django app

Execution Environment

- OpInt Framework has been deployed on CERN OpenShift and CERN OpenStack VMs
- Resources provisioned by us for the web application are optimized for ML.
- ML needs to be delegated to dedicated ML clusters or better K8s based backend.
- We used workers in the CERN Spark clusters for ML tasks.
 - OpenStack VM acts as a spark submit node
 - Runs OpInt Framework + Luca's implementation of NLP adapters.
 - Spark_submit with command line options used to execute the ML jobs.
- Panos is looking into K8s based deployments

What's Next

- Look into dedicated resources for ML workloads.
 - Spark team would be able to provision a spark queue for this purpose.
 - GPUs?
 - Interface OpInt Framework to interact with these resources.
-
- Evaluation of results is required
 - How to convince people to invest time into it?
 - Cross check with GGUS tickets as pre-validation to show the approach has potential
 - Improve result visualization (Make it more user friendly and intuitive)
 - Involve Experts
-
- Result exposure to [live page](#) (ADCoS)
 - [Annotations](#) of results?