Intelligent Anomaly Detection in High Performance Computing Logs via Machine Learning

Alexandra DeLucia (Rollins College)
Mentor: Lissa Baseman, HPC-DES

USRC
Ultrascale Systems Research Center (ultrascale.org)

Los Alamos NATIONAL LABORATORY EST.1943

Motivation

- **Goal:** Predict job outcomes from system log (syslog) messages
- The reason for an incomplete job outcome is frequently hidden within many lines of corresponding syslog
- Current state-of-the-art is for human operators to sift through both joblog and syslog datasets by hand
- Each LANL HPC machine stores each node's syslog as well as the job scheduler's joblog, separately

Experimental Setup

- Gather syslog and joblog files from Wolf, a LANL HPC cluster, from the same time frame
- For each job, extract the corresponding syslog chunk from the relevant time span and nodes
- Extract relevant features from job-specific syslog chunks
- Measure predictive performance on job outcome using two supervised classification algorithms, varying the input features

Feature Extraction

- Automatically extract topics and topic distributions on each job-specific syslog chunk using latent Dirichlet allocation (LDA) [1], as implemented in MALLET [2]

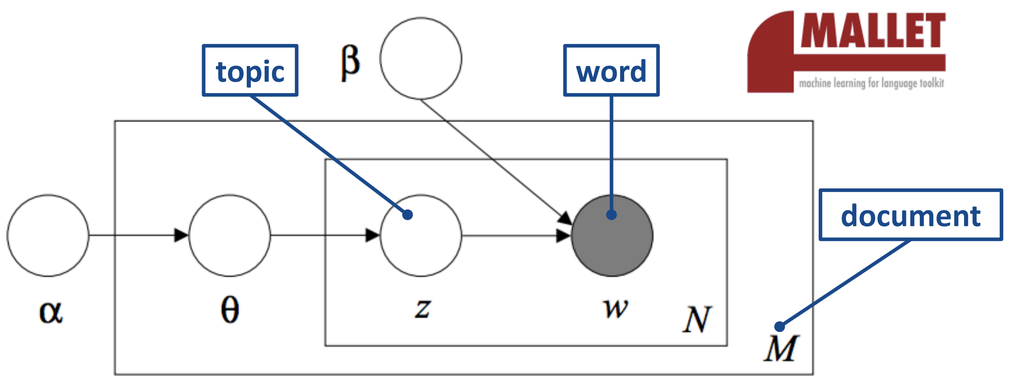

Graphical model representation of LDA, from [1].

Topic	Keywords
1	session user pam unix sshd uid closed opened access granted root
2	iterations msgsize ranks pattern jobid wolf cluster sequential avgbw tonode fromnode send read usec avglat test write session aggregate ring
3	write sys file error channel socket addr cpu memory not page open task obdclass time own read check generic area
4	session sshd root pam unix user isa coretemp physical closed opened uid
5	session lustre mdt user root mdc sshd connection pam unix init fault obdclass coretemp physical isa uid will lost operations
6	session unix user pam root sshd opened uid closed
7	session pam sshd user unix root physical coretemp isa lnet connection lustreerror import closed inet opened uid own server routers
8	ost lustre osc connection previous similar skipped messages restored progress will service operations using lost recovery complete message kiblnd iblnd

- Include other basic features, such as average time between syslog messages, and total time difference between first and last messages

Results

- Topic distributions for each job state were unique

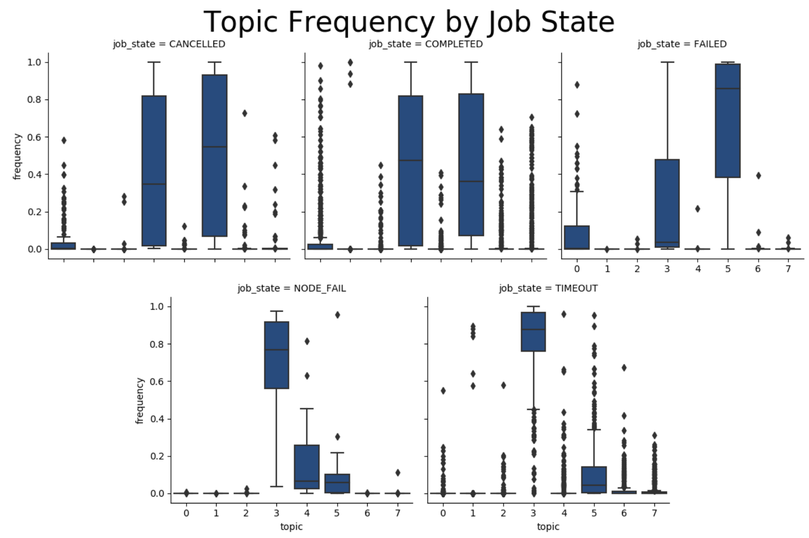

- Random forests had greater ROC AUC scores than support vector machines in joblog classification across the number of topics

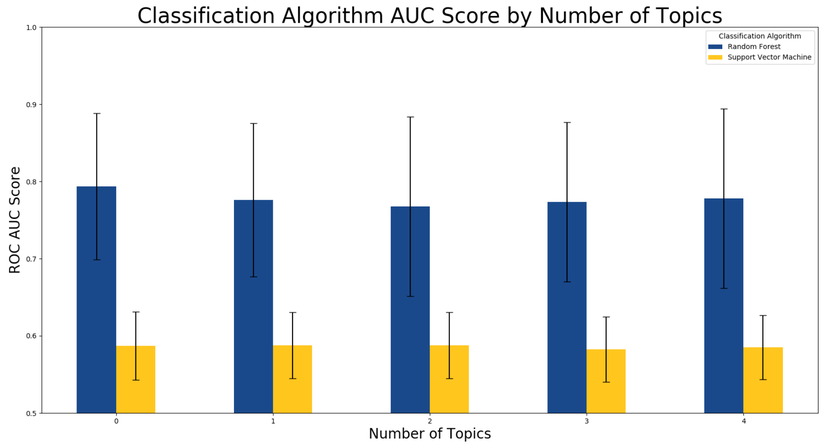

Conclusion

- Basic features and topics from syslog predict job outcome with moderate success, particularly with random forests
- Extracted topics show significant signal across job outcome states, but adding additional intelligent features may improve performance

Summary

- Human operators must currently determine reasons for undesirable job outcomes by hand using system logs (syslog)
- Supervised machine learning techniques can successfully predict job outcomes from corresponding syslog messages, significantly reducing the required human effort

[1] Blei, David M., Andrew Y. Ng, and Michael I. Jordan. "Latent Dirichlet Allocation." JMLR, 2003.
[2] MaCallum, Andrew Kachites. "MALLET: A Machine Learning for Language ToolKit." http://mallet.cs.umass.edu. 2002.

LA-UR-TODO