LOG-BASED FINGERPRINTING OF HPC APPLICATIONS

Ryan D. Lewis[1], Zhengchung Liu[2], Rajkumar Kettimuthu[2], Michael E. Papka[1,3]

1 Department of Computer Science Northern Illinois University
2 Data Sciences and Learning, Argonne National Laboratory
3 Argonne Leadership Computing Facility, Argonne National Laboratory

INTRODUCTION AND MOTIVATION

- High Performance Computing (HPC) is an important method for scientific discovery via large scale simulation, data analysis, or artificial intelligence.
- In order to improve our understanding of HPC applications, user demands, and resource usage characteristics at Leadership Computing Facilities, we perform correlative analysis of various logs for different subsystems of ALCF's Mira leadership supercomputer.
- Based on insights gained from our analysis, we engineer features to "fingerprint" an HPC application.
- Our initial analysis shows that the largest consumers of core-hours, both users and applications, exhibit predictable patterns of usage, making them promising targets for classification and prediction efforts.

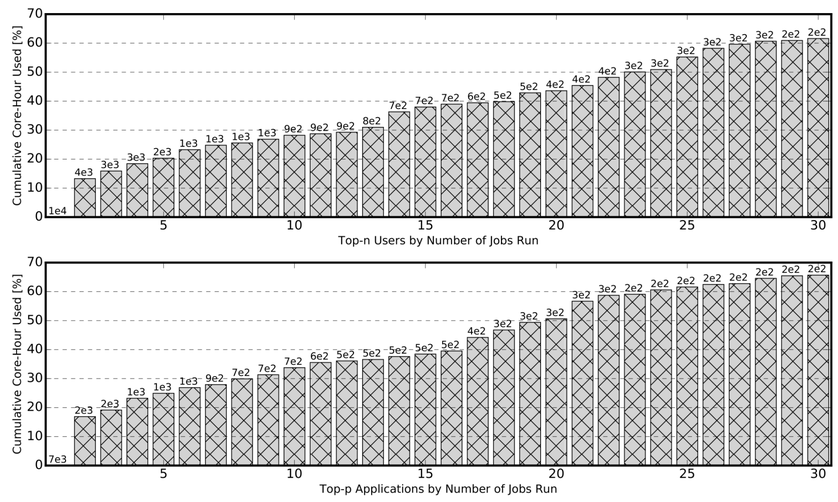

VERIFYING APPLICATION IDENTITY

- For auditing and security purposes, it is useful to ensure that jobs using HPC resources are correctly identifiable.
- t-SNE representations allow us to visualize complex higher dimensional representations in 2D form, where we can see clear clustering of different executables.
- With this in mind, we were able to construct machine learning algorithms using XGBoost and Neural Network that could accurately predict the application with >99% accuracy.

- Features Engineered:
  - T_{comm}: Fraction of runtime spent on inter-process communication.
  - T_{fio}: Fraction of runtime on file IO.
  - OPS: Achieved operations per second.
  - FLOPS: Achieved floating point operations per second per node.
  - PPN: Number of processes per node.
  - MEM_F: Average RAM fetch per cycle
  - MEM_S: Average RAM store per cycle

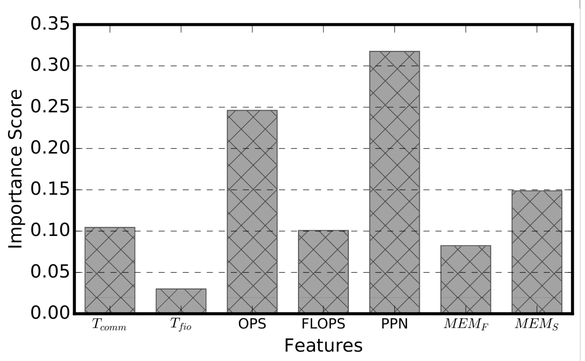

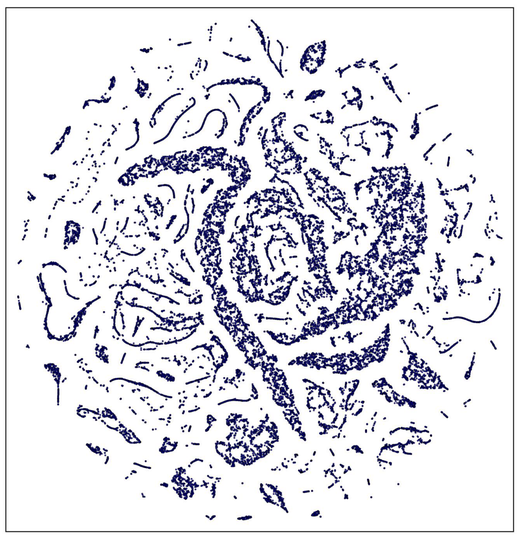

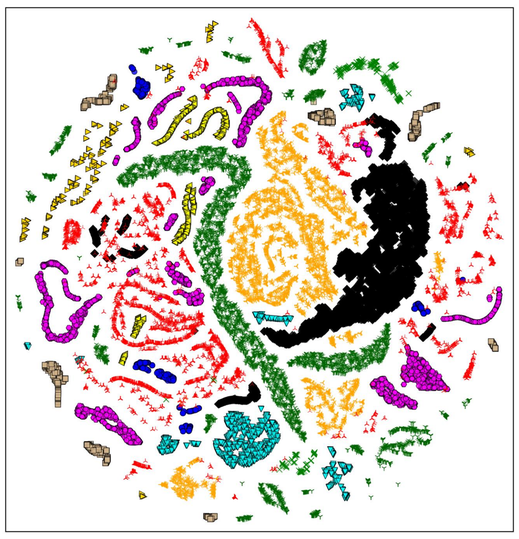

2D t-SNE embeddings of the task representation based on the 7 training features described above. Color encodes the application name.

PREDICTING APPLICATION RUNTIME

- More accurate runtime prediction of jobs has major benefits for job scheduling on leadership computers but is challenging due to limited information at job submission.
- We attempt to take advantage of the repetitive behavior of top core-hours consumers on Mira to train an accurate machine learning model.
- We trained our model over the first 9.6 months (80%) of a year's worth of data using only information available at submission time, and the remaining 20% to test.

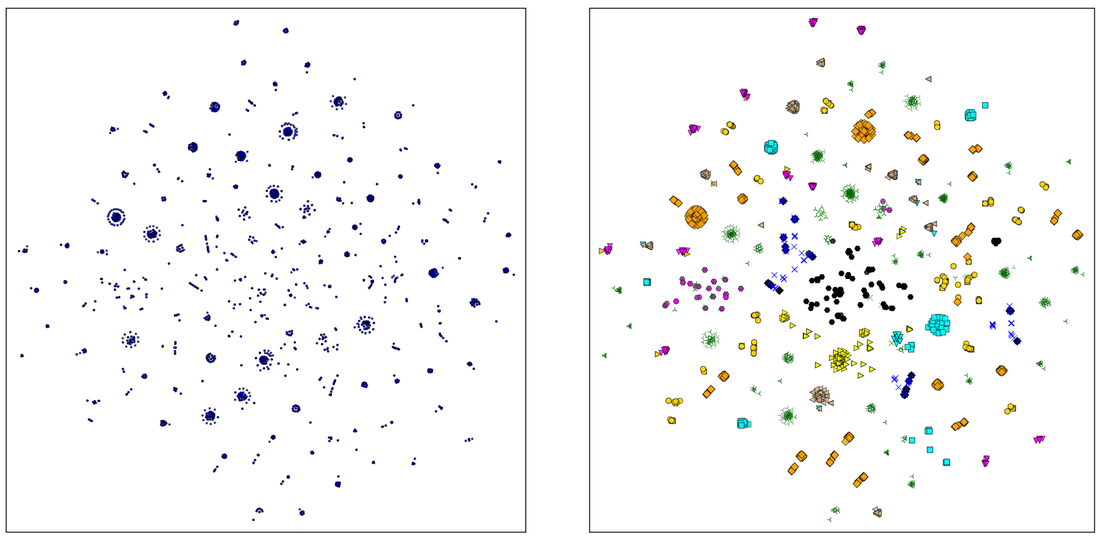

2D t-SNE embeddings of the task representation based on information available at job submission. Color encodes the applications.

- Due to limited information at submission time, perfectly accurate predictions are difficult.
- Poor predictions can lead to unacceptable outcomes, such as prematurely killed jobs.

Margin time (min)	Underestimated jobs	Beneficial Jobs
10	424	2,620
30	314	1538
60	150	582
90	73	493
180	7	412
310	0	30

- Potential Mitigation: add a fixed grace period to predictions, to reduce the likelihood of premature job killings.

IMPACT

- The ability to verify the identity of applications running on leadership supercomputers helps ensure that scientific resources are being used appropriately and as intended.
- Improved runtime prediction can lead to more efficient scheduling, allowing for improved response time.

FUTURE WORK

- The performance logs used in this project can provide many features for fingerprinting and classifying applications, but the collectors also introduce overhead, sometimes considerable, to the application.
- To avoid this, future work may attempt to build machine learning based classifiers, trained with labels and hardware performance counters as input to reduce the coverage of jobs (e.g. randomly select some jobs) that must collect performance logs.
- Another potential area of future development combines our current models with a tracked libraries dataset, allowing us to do analysis on and make predictions with the libraries applications are using, for better pre-runtime prediction.

U.S. DEPARTMENT OF ENERGY
Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

Northern Illinois University
Argonne NATIONAL LABORATORY