

# Web Application For Identifying and Diagnosing Performance Anomalies

Kyra Terez<sup>1,2</sup>, Efe Sencan<sup>2</sup>, Prof. Ayse K. Coskun<sup>2</sup>

<sup>1</sup>Travis High School, 11111 Harlem Road, Richmond, TX 774061;

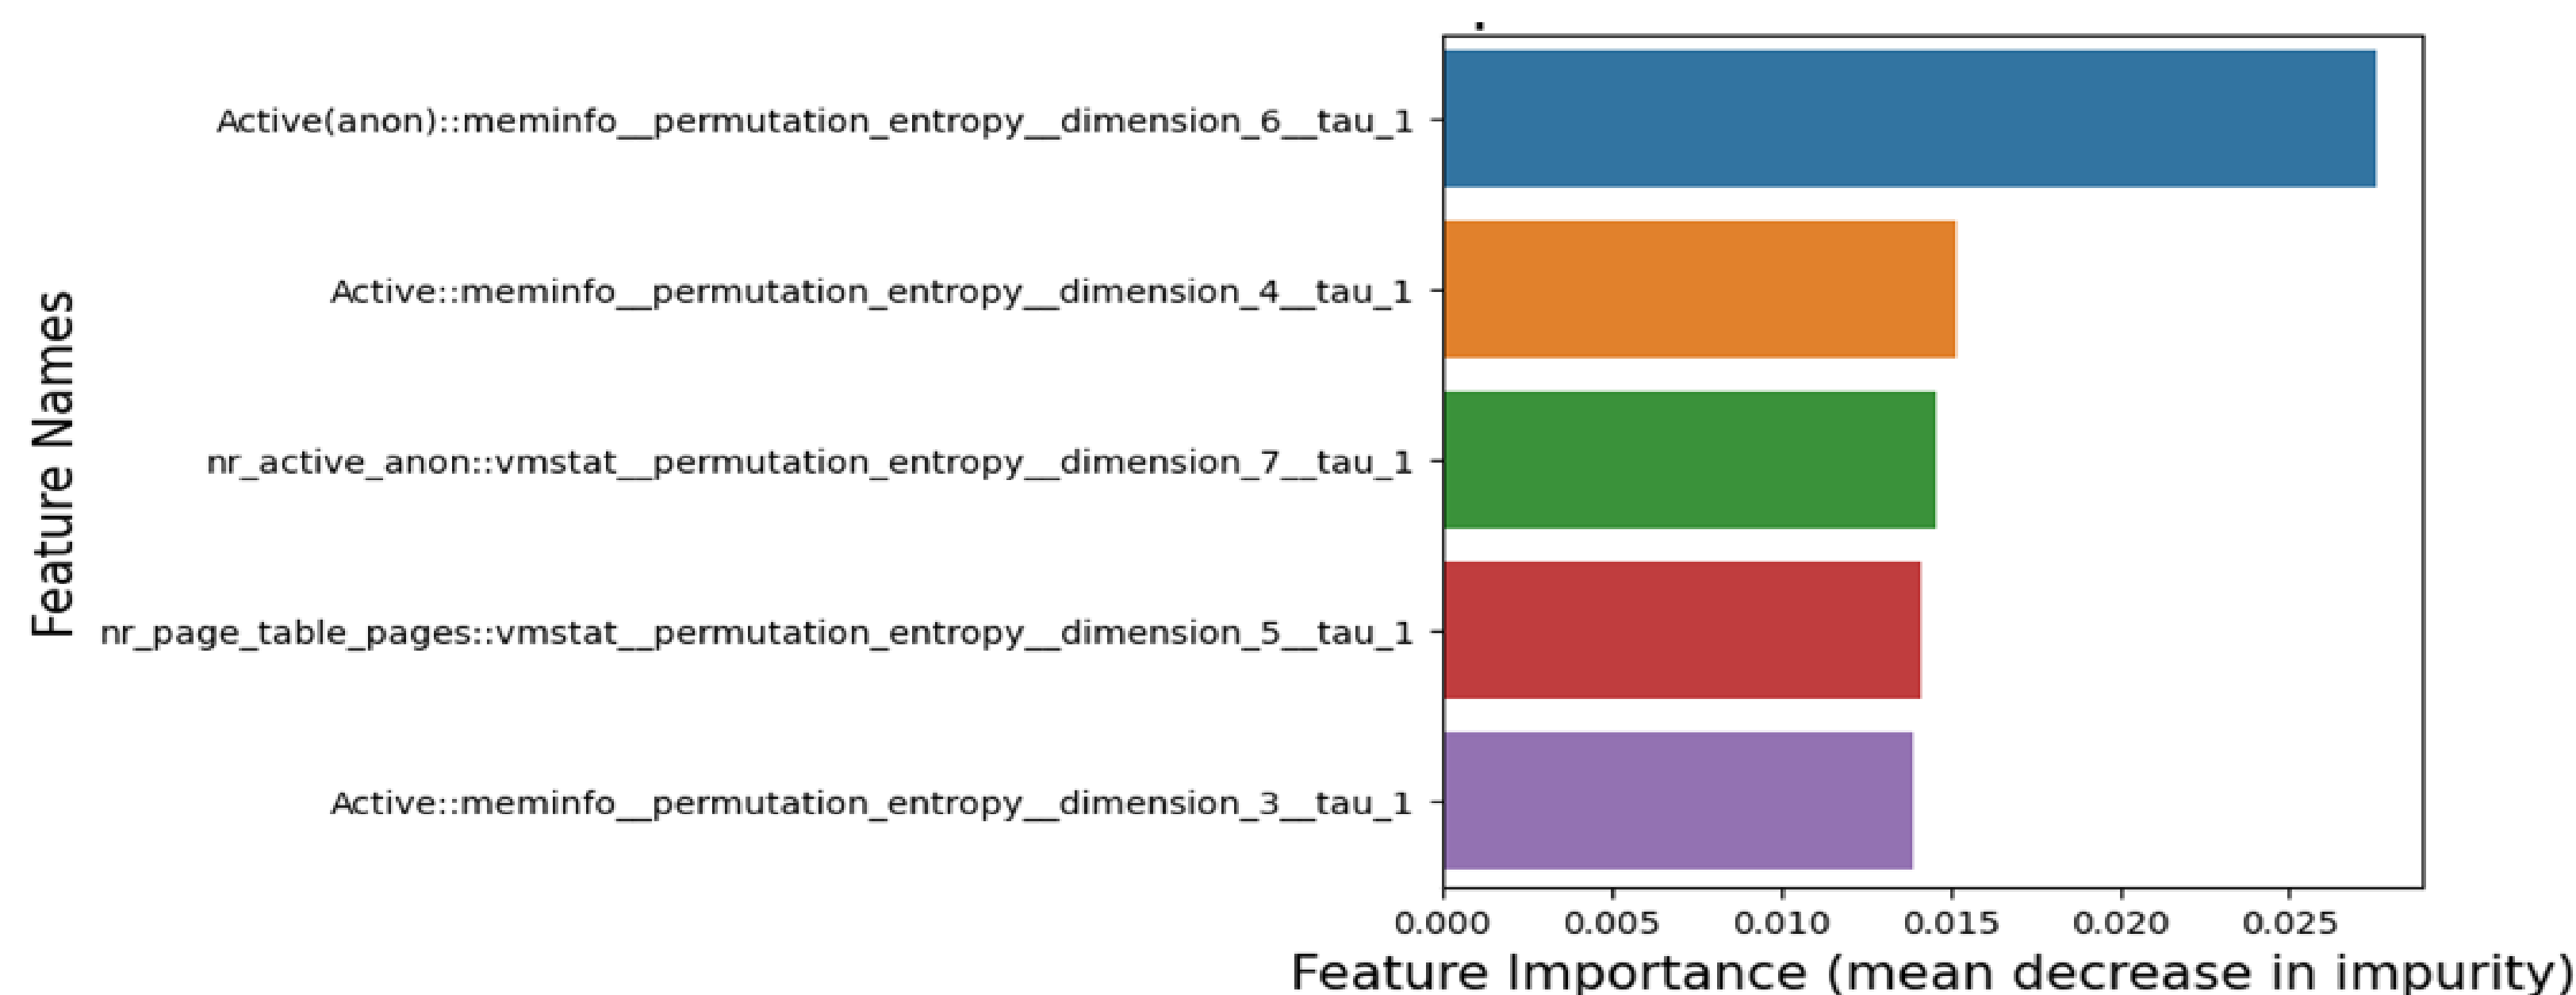
<sup>2</sup>Department of Electrical and Computer Engineering, Boston University, 8 St. Mary's Street, Boston, MA 02215

## Introduction

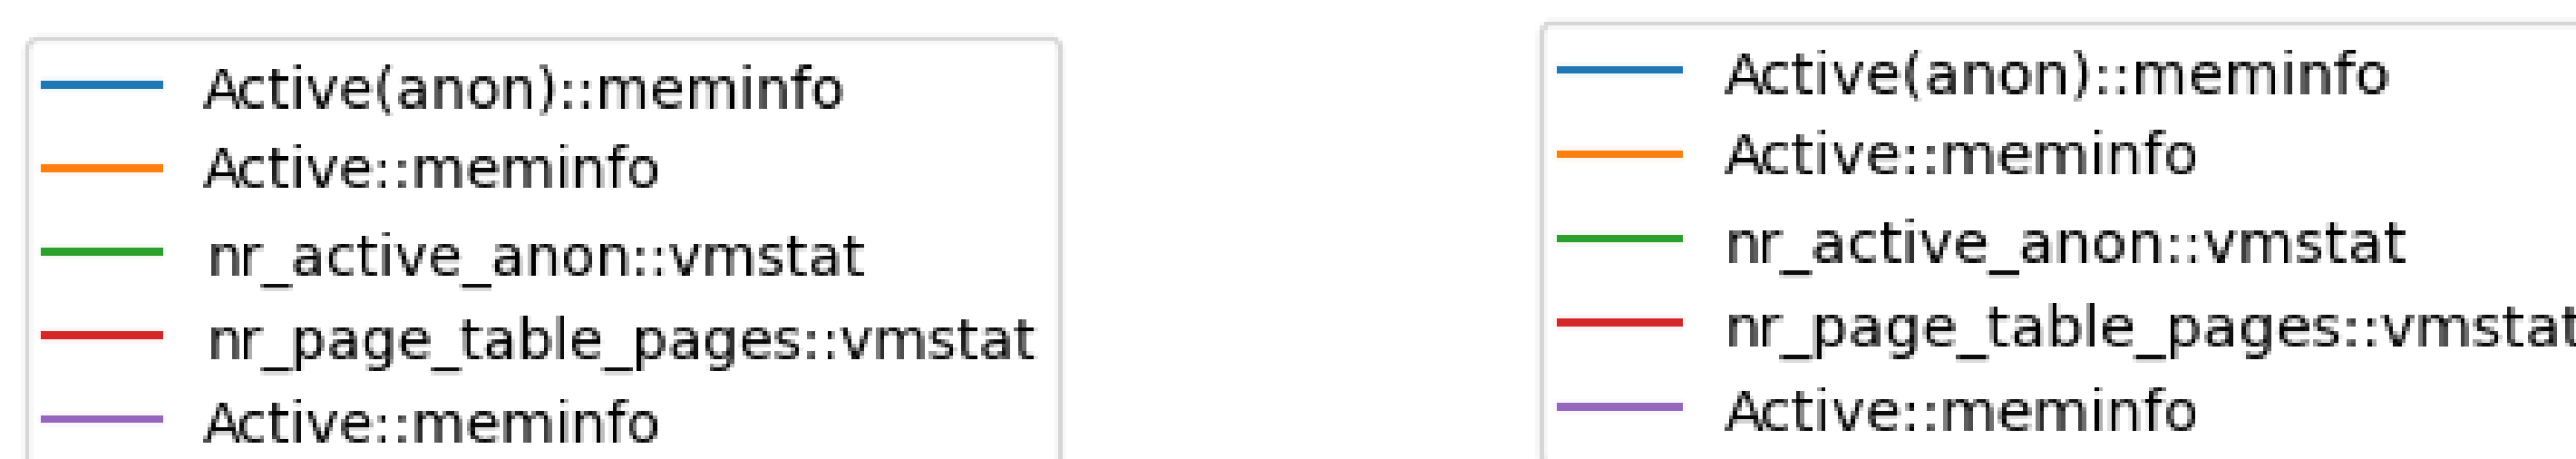
- High Performance Computing (HPC) Systems encounter performance variations
- Machine learning frameworks have been developed
  - Detects performance anomalies on compute nodes
  - Complex for HPC system users
- Created a web application for users using a machine learning framework
  - Flask micro-web framework
  - Gives HPC system users results from detecting performance anomalies
  - Using anomaly types: memeater, dcopy, leak, dial

## Results


The application displays overall information about the application run and drill-down application run analysis



Top feature importance scores from model training displayed in overall results

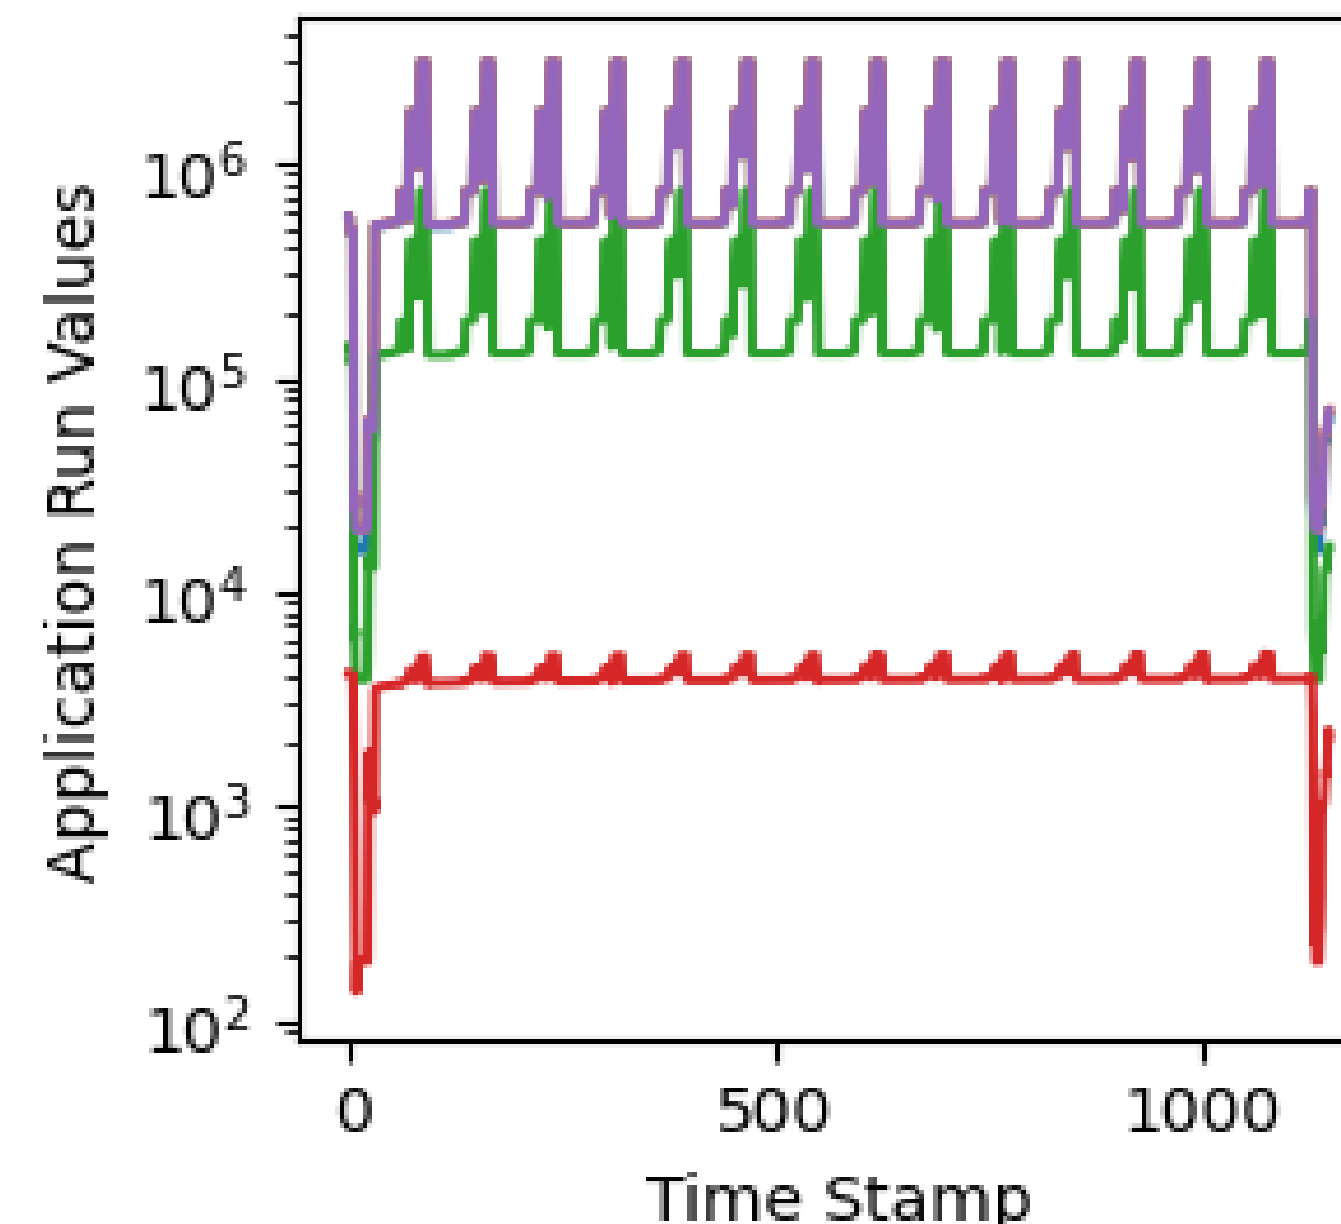


Top Metrics for Healthy Application Run



Telemetry data for sample healthy application run based on the feature importance scores

Top Metrics for Application Run Selected

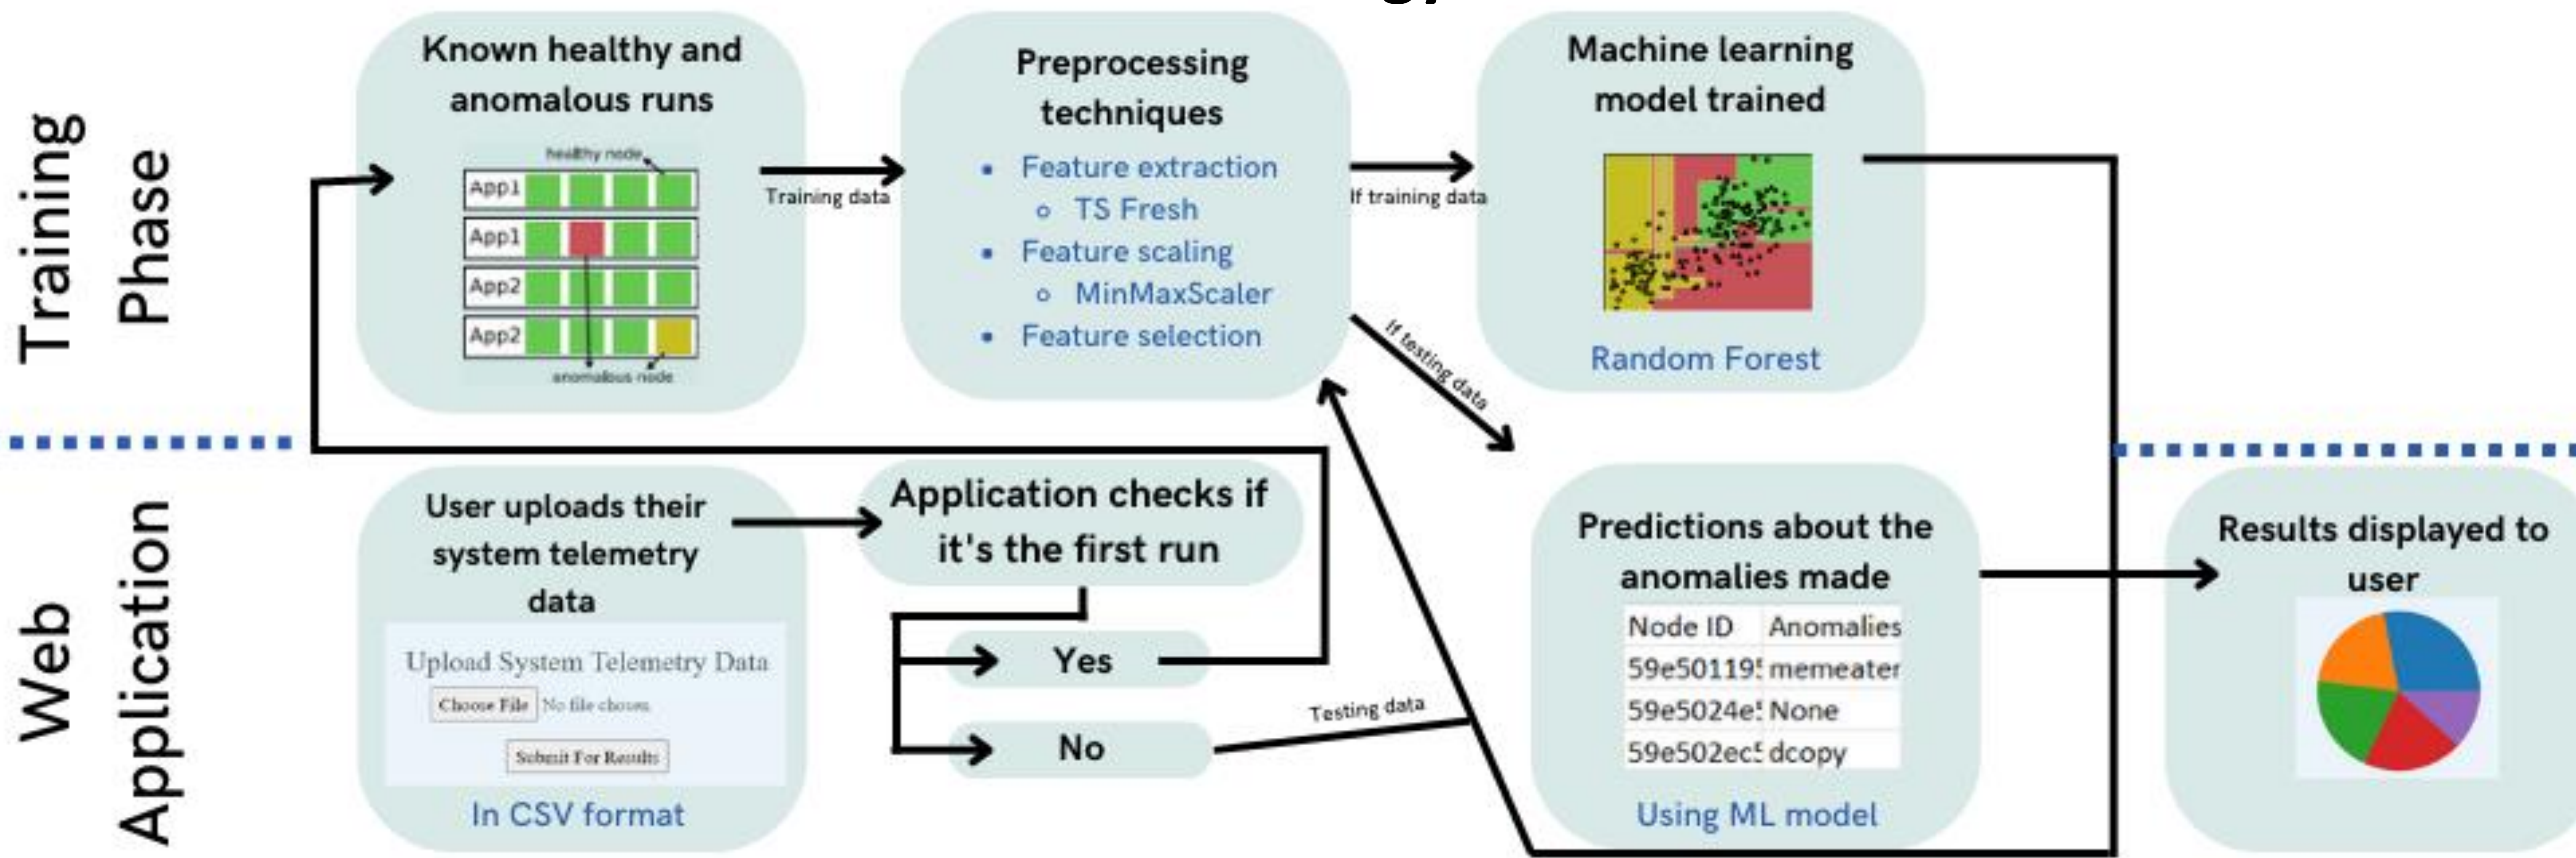


Telemetry data for specific application run with dcopy anomaly type

## Discussion/Conclusions

- HPC system users should understand the compute node anomalies within their system
- This allows them to make system adjustments to account for performance variations
- The web application creates different components that allow users to obtain this understanding
  - Users can see the feature importance scores
  - Users can compare the sample healthy application run's telemetry data to the anomalous run's data

## Methodology



## References

[1] Tuncer, O.; Ates, E.; Zhang, Y.; Turk, A.; Brandt, J.; Leung, V. J.; Egele, M.; Coskun, A. K. Online Diagnosis of Performance Variation in HPC Systems Using Machine Learning. *IEEE Transactions on Parallel and Distributed Systems* 2019, 30 (4), 883–896.

## Acknowledgements

Thank you to Efe Sencan and Professor Coskun for the guidance and help given throughout this project. I would also like to thank my research partner, Richard Chen, and the RISE program for this opportunity