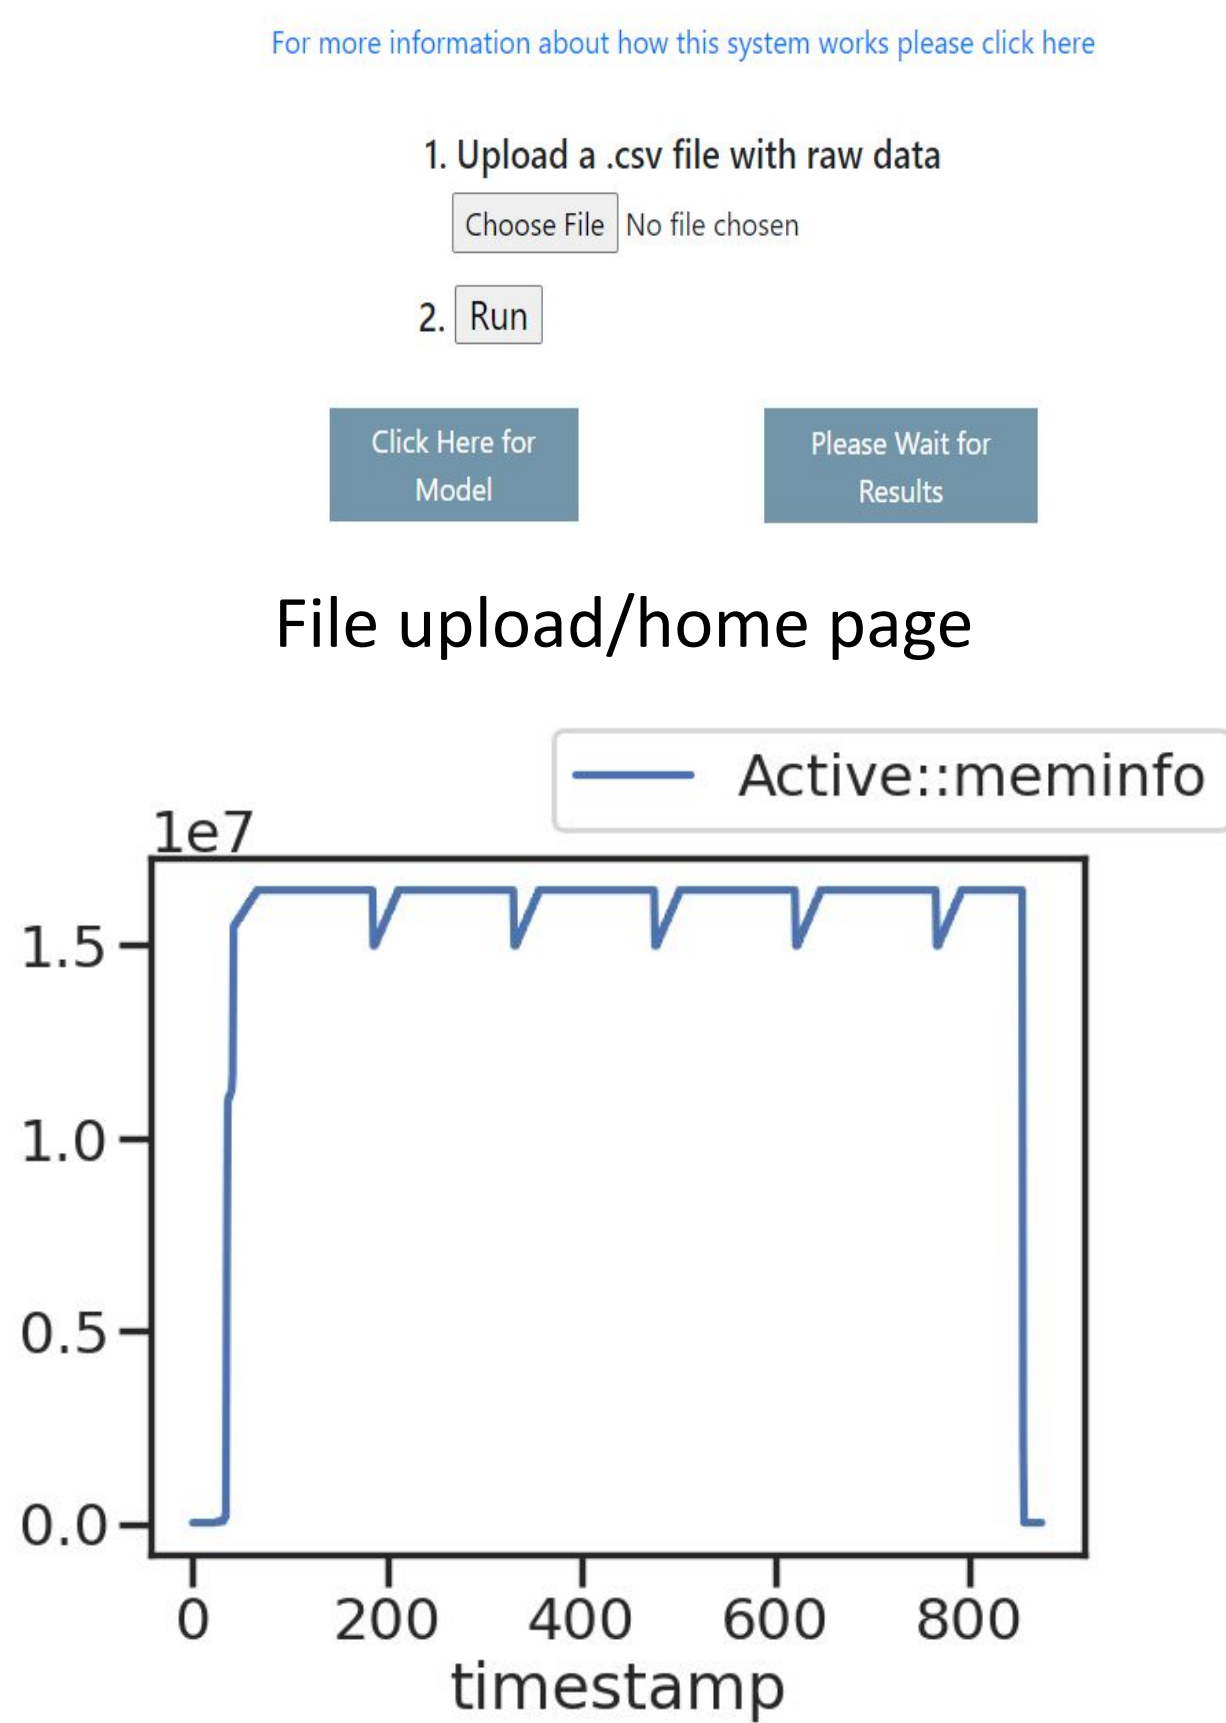


Introduction

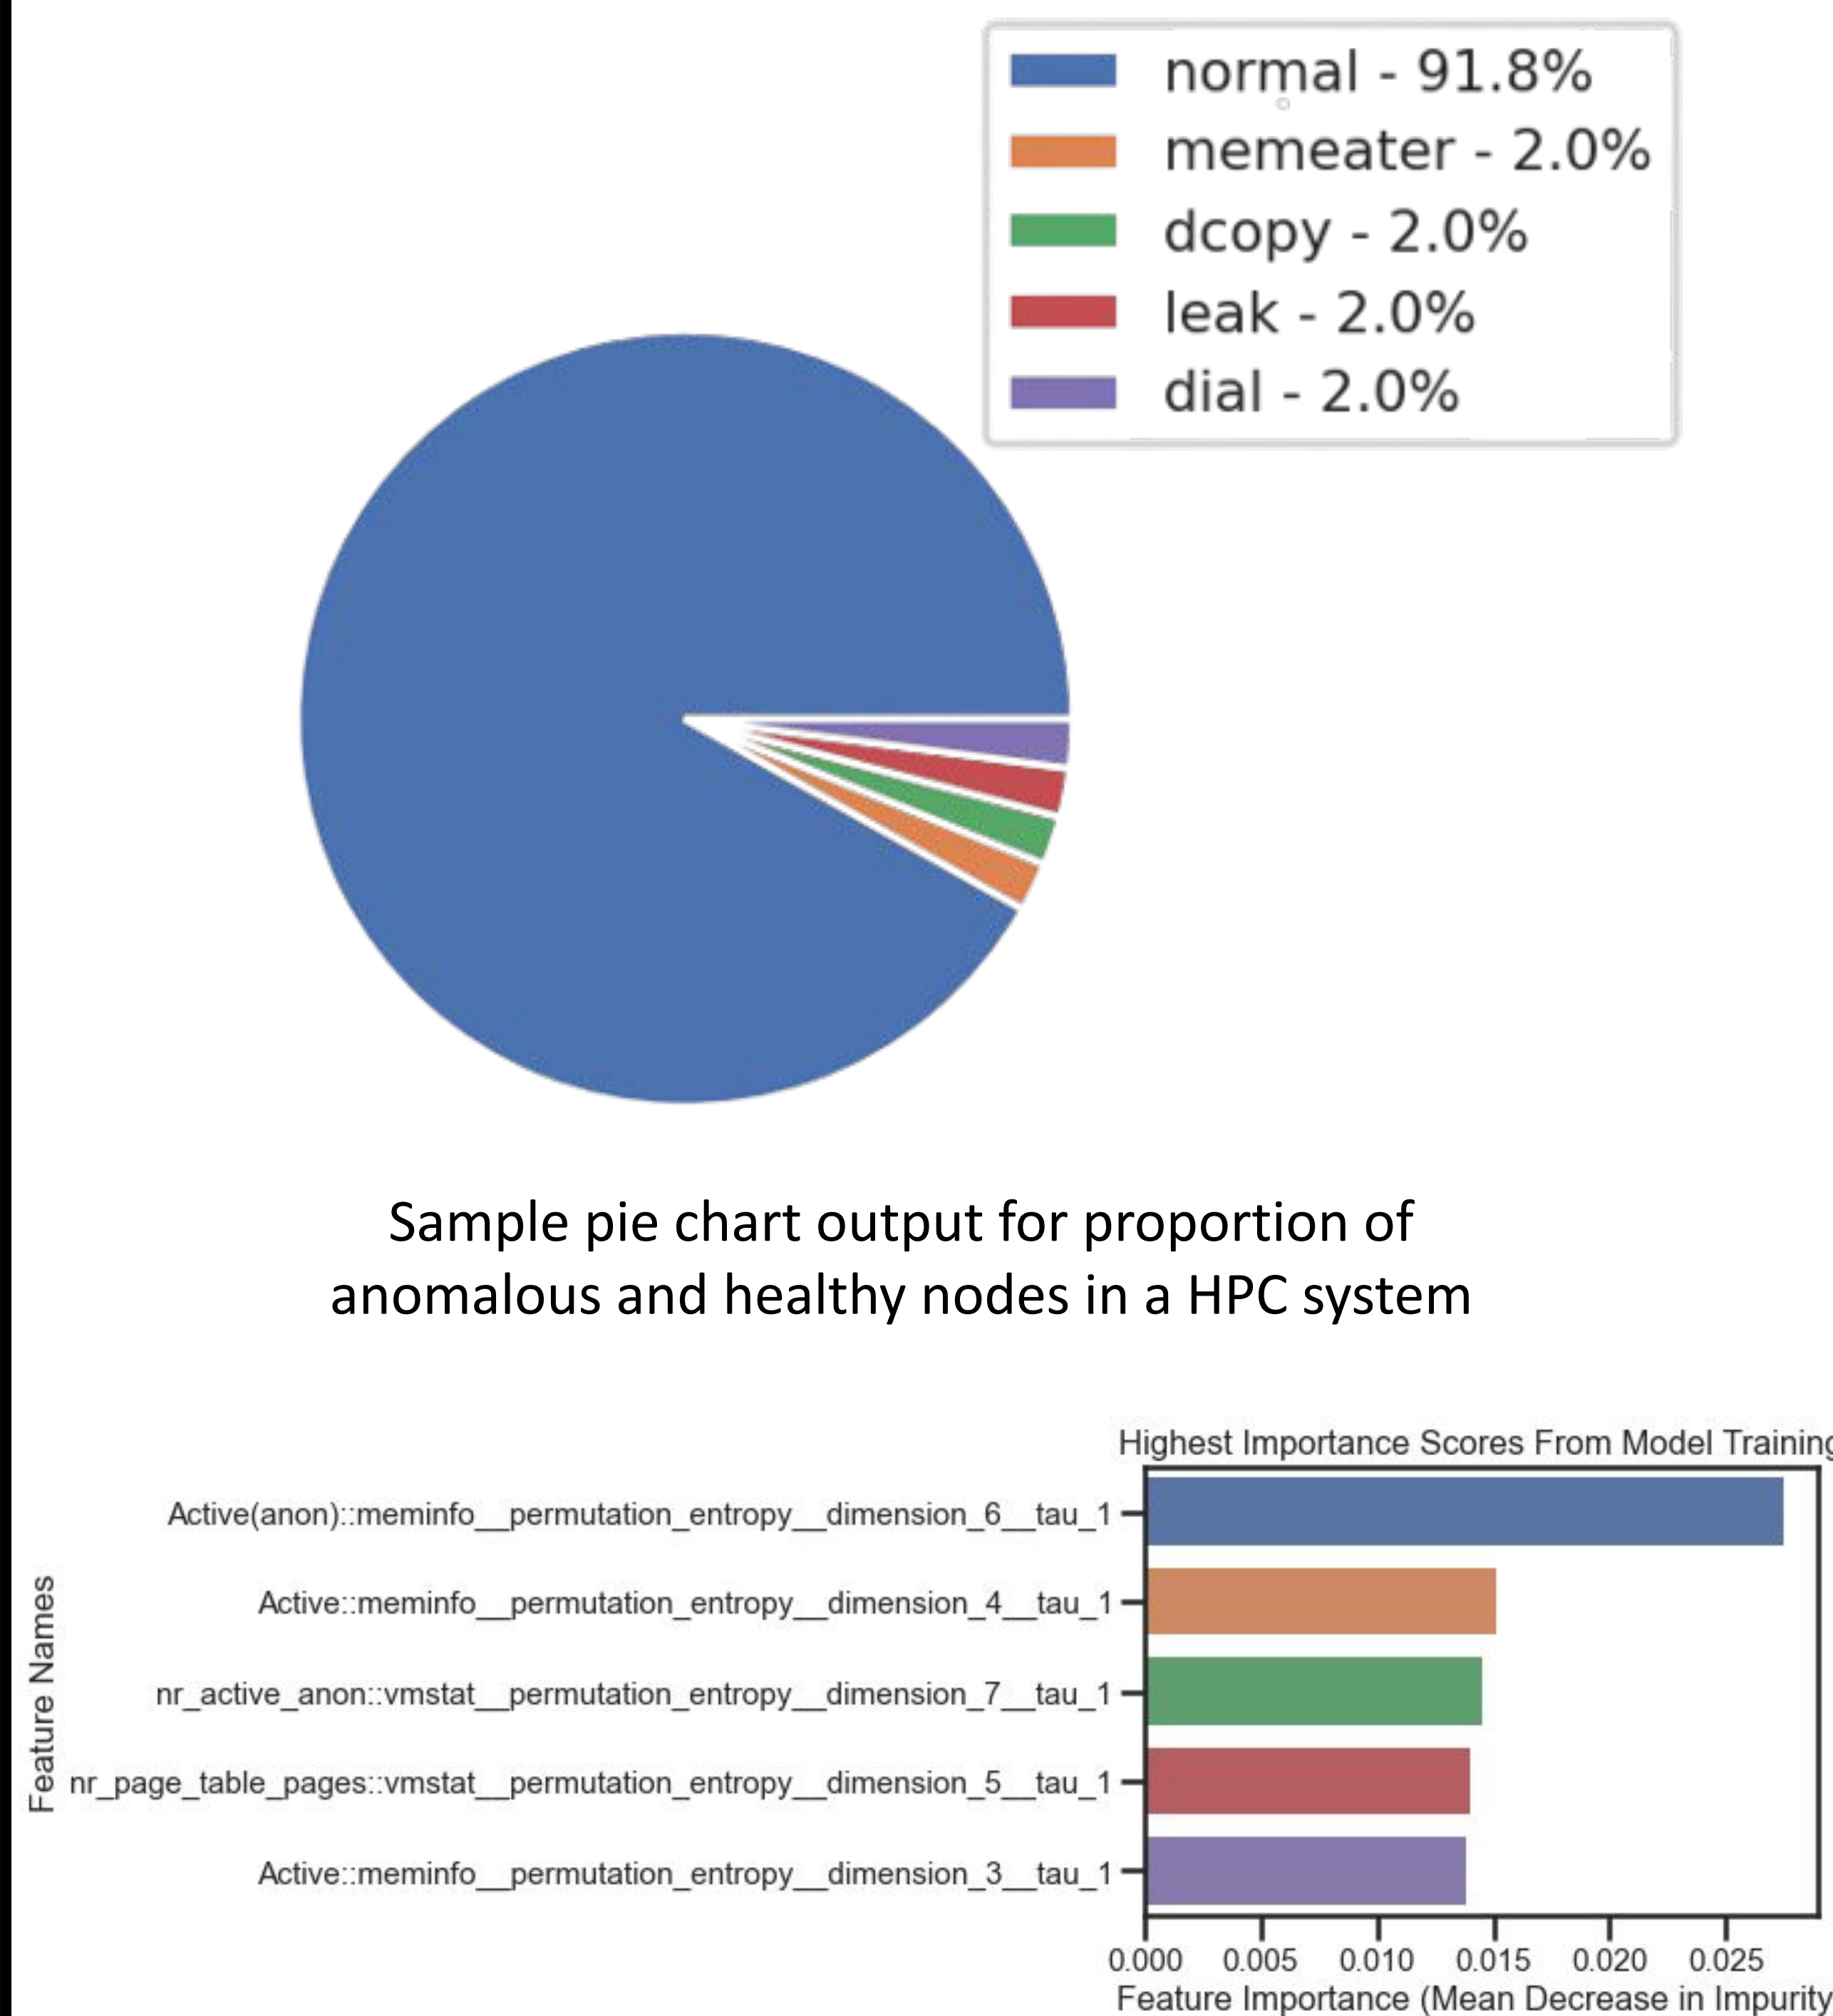
- High-Performance Computing (HPC) systems growing larger and more complex
 - Runs scientific simulations and AI
- HPC Systems suffer from performance variations
 - Resource contention, memory leaks, and hardware related problems
 - Leads to higher energy consumption
- Machine Learning (ML) models created
 - Hard to use
 - Outputs are complicated to interpret for lay people
- Synthetic anomalies diagnosed in model:
 - dcopy
 - memeater
 - dial
 - leak
- Web application created in order to produce better insight

HPCWeb: A web based HPC diagnostic system

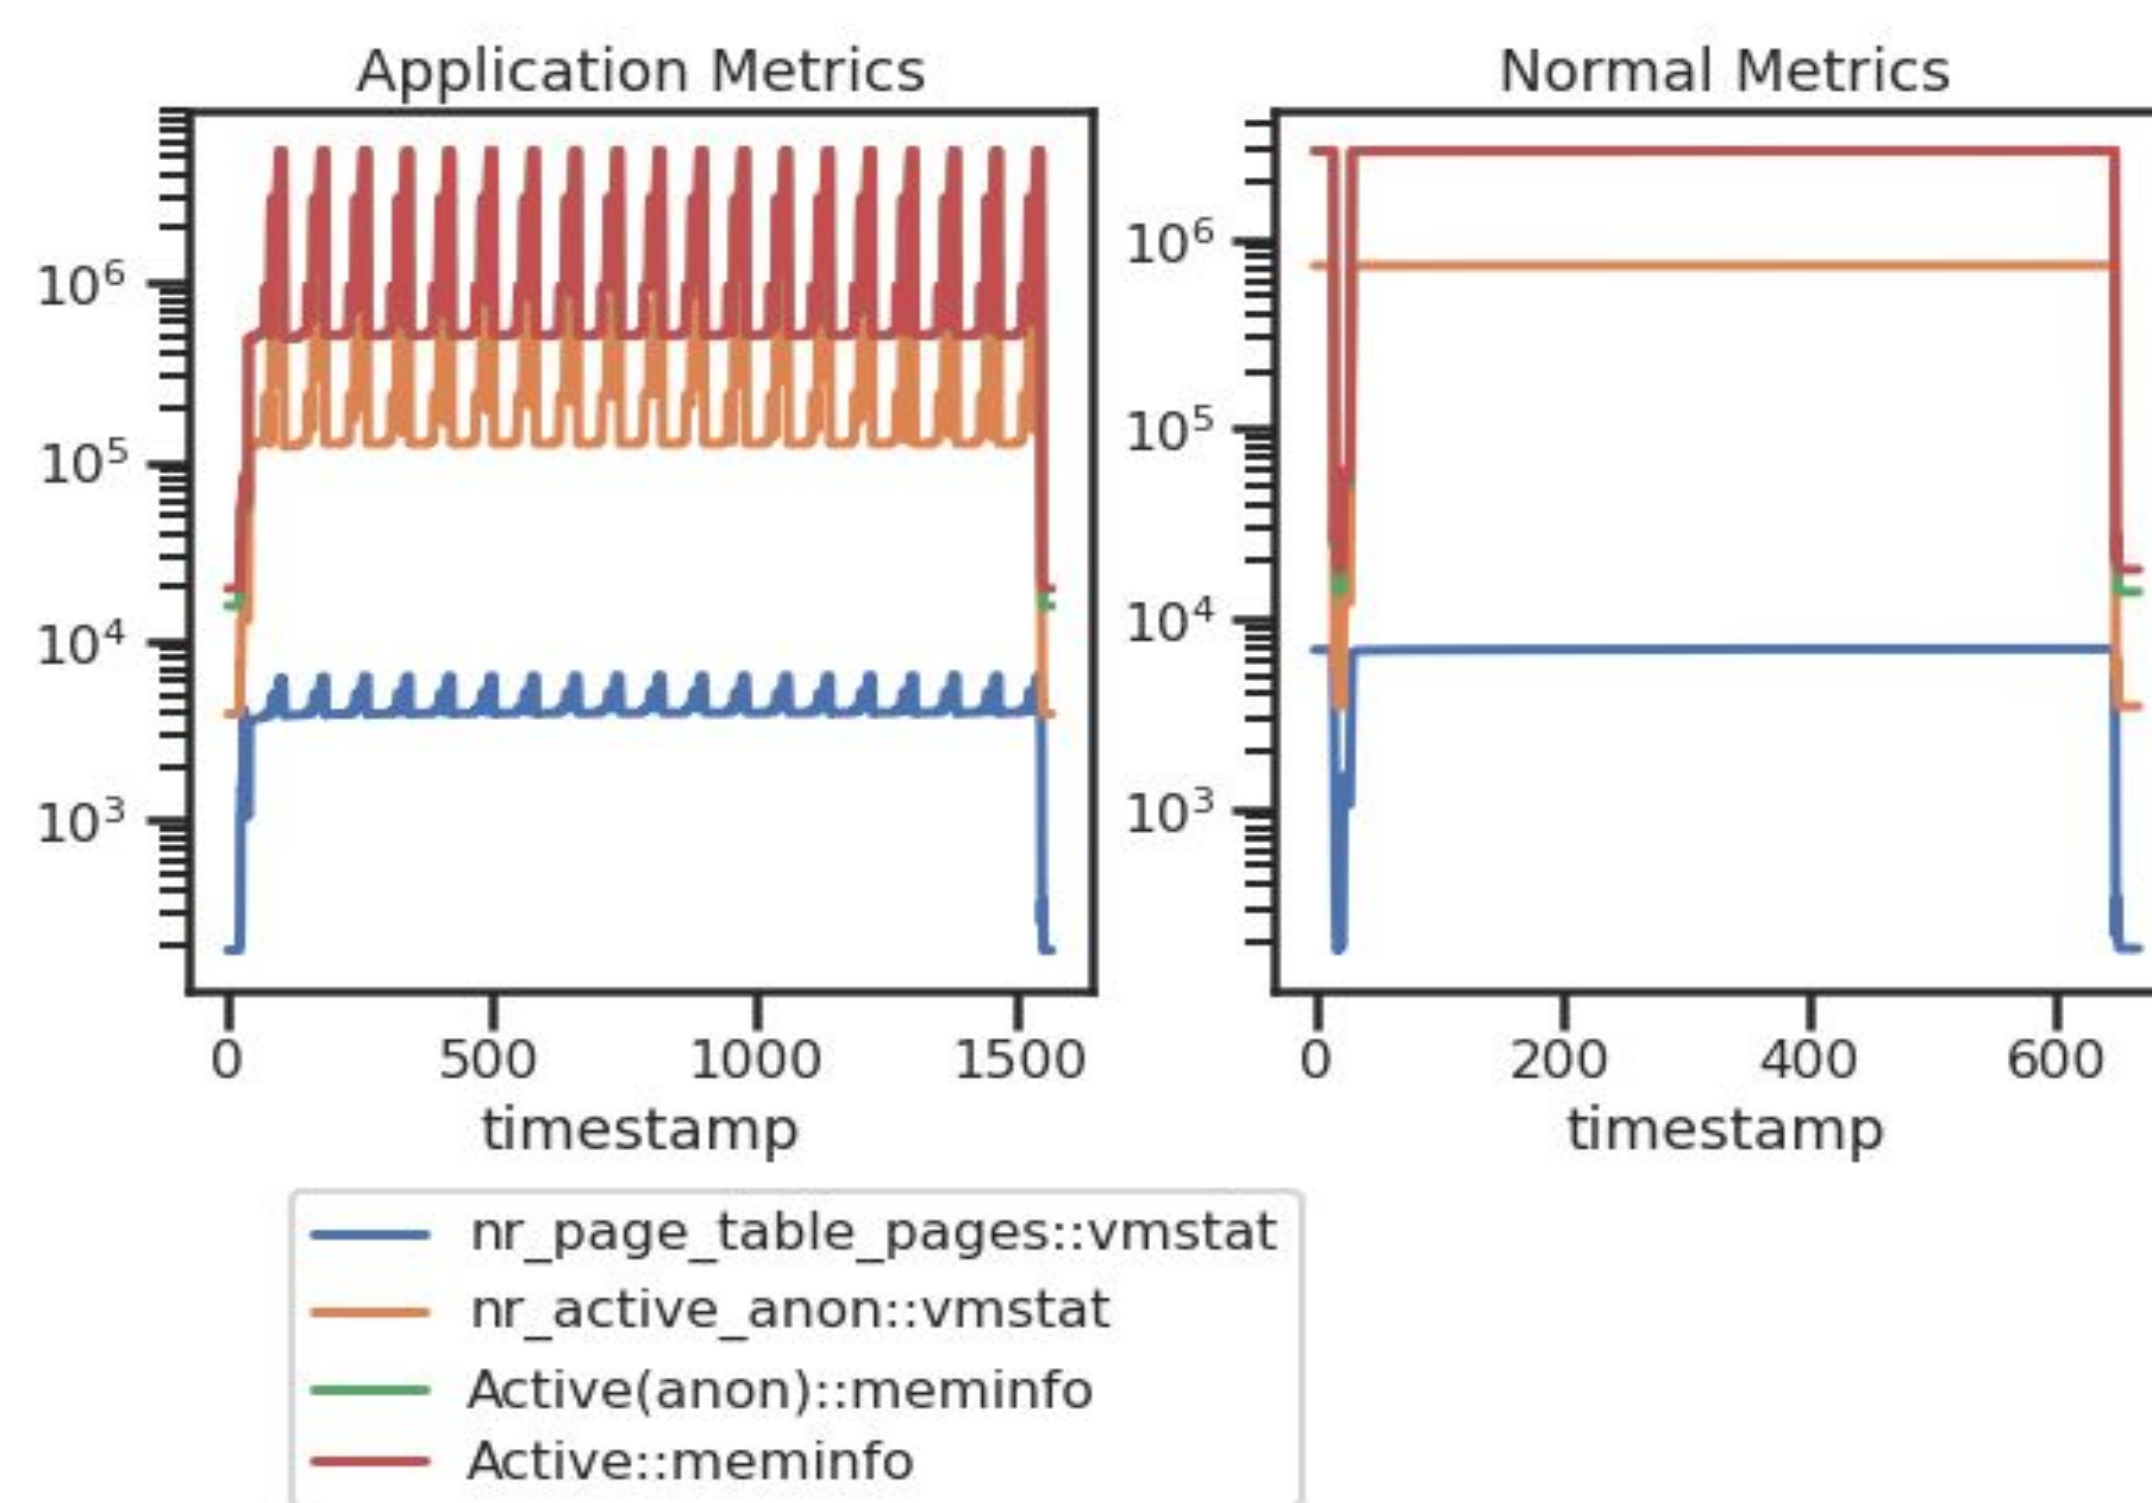
Instructions
Please upload a csv file to be predicted using the machine learning model. Once uploaded, you can find out the predictions that were made about the anomaly types (if found) on the results page.



Results



Most important features during model training



Example comparison between an anomalous process (left) and a healthy process (right)


Discussion/Conclusions

- HPC systems increasing in computational power
 - Performance variations more likely
 - Easy to use diagnosis frameworks more important
- Anomalies have different metric signatures
- Real world scenarios consist mainly of healthy job runs
- Easily gain insight into HPC systems and performance variations within the system
 - HPC users know what led to diagnosis
 - Users can see visuals of the ratio of anomalies
 - Side by side healthy vs anomalous graphs
 - Easily see distinctions
 - Distinguish what features led to anomaly diagnosis

Future Work

- Integrate methods for runtime diagnosis
 - Allows users to see exact time an anomaly occurs within system
- Deploy website onto a web server in future
 - Allows more people in the HPC community to gain access to ML model
 - Get user feedback to improve functionality of website
- Improve efficiency

Methodology



1. User interface created using Python Flask and HTML
2. ML model trained using train data and labels
3. Efficacy of model tested using raw telemetry data
4. Predicted results shown in charts developed using Matplotlib

References

- [1] Aksar, B., Schwaller, B., Aaziz, O., Leung, V. J., Brandt, J., Egele, M., & Coskun, A. K. (2021). E2EWatch: An end-to-end anomaly diagnosis framework for production HPC Systems. *Euro-Par 2021: Parallel Processing*, 70–85. https://doi.org/10.1007/978-3-030-85665-6_5
- [2] Tuncer, O., Ates, E., Zhang, Y., Turk, A., Brandt, J., Leung, V. J., Egele, M., & Coskun, A. K. (2019). Online diagnosis of performance variation in HPC systems using Machine Learning. *IEEE Transactions on Parallel and Distributed Systems*, 30(4), 883–896. <https://doi.org/10.1109/tpds.2018.2870403>

Acknowledgements

I would like to thank Efe Sencan and Professor Coskun for supporting and mentoring me throughout the entire project. I would also like to thank the rest of PEACLab for providing me with help along the way. Lastly, I would like to thank the RISE program for providing me with the amazing experience to learn and conduct research in a formal setting.