

AI4HPC: AI-based Scalable Analytics for Improving Performance, Resilience, and Security of HPC Systems

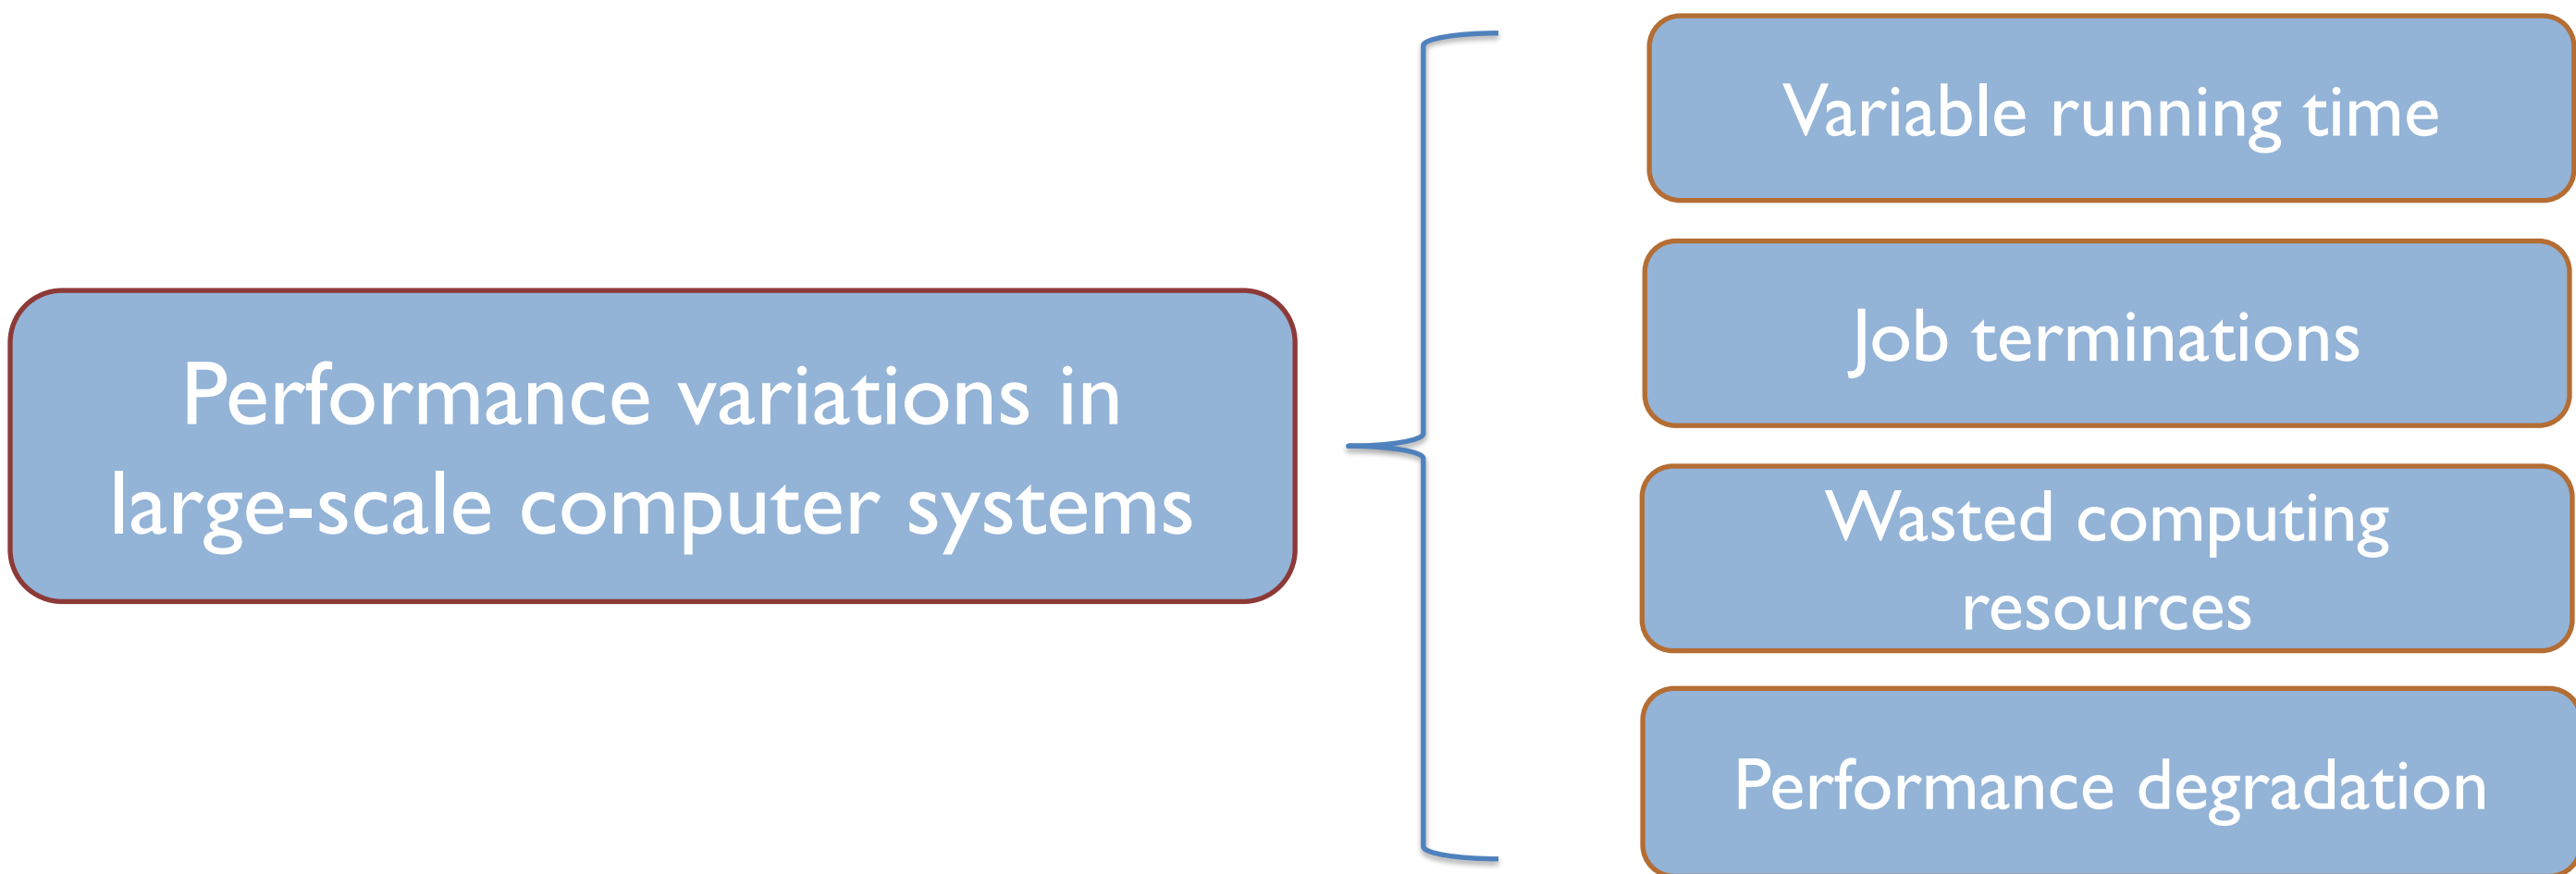


Efe Sencan, Yin-Ching Lee, Burak Aksar, Richard Chen, Ayse K. Coskun
For contact: leewill@bu.edu, esencan@bu.edu



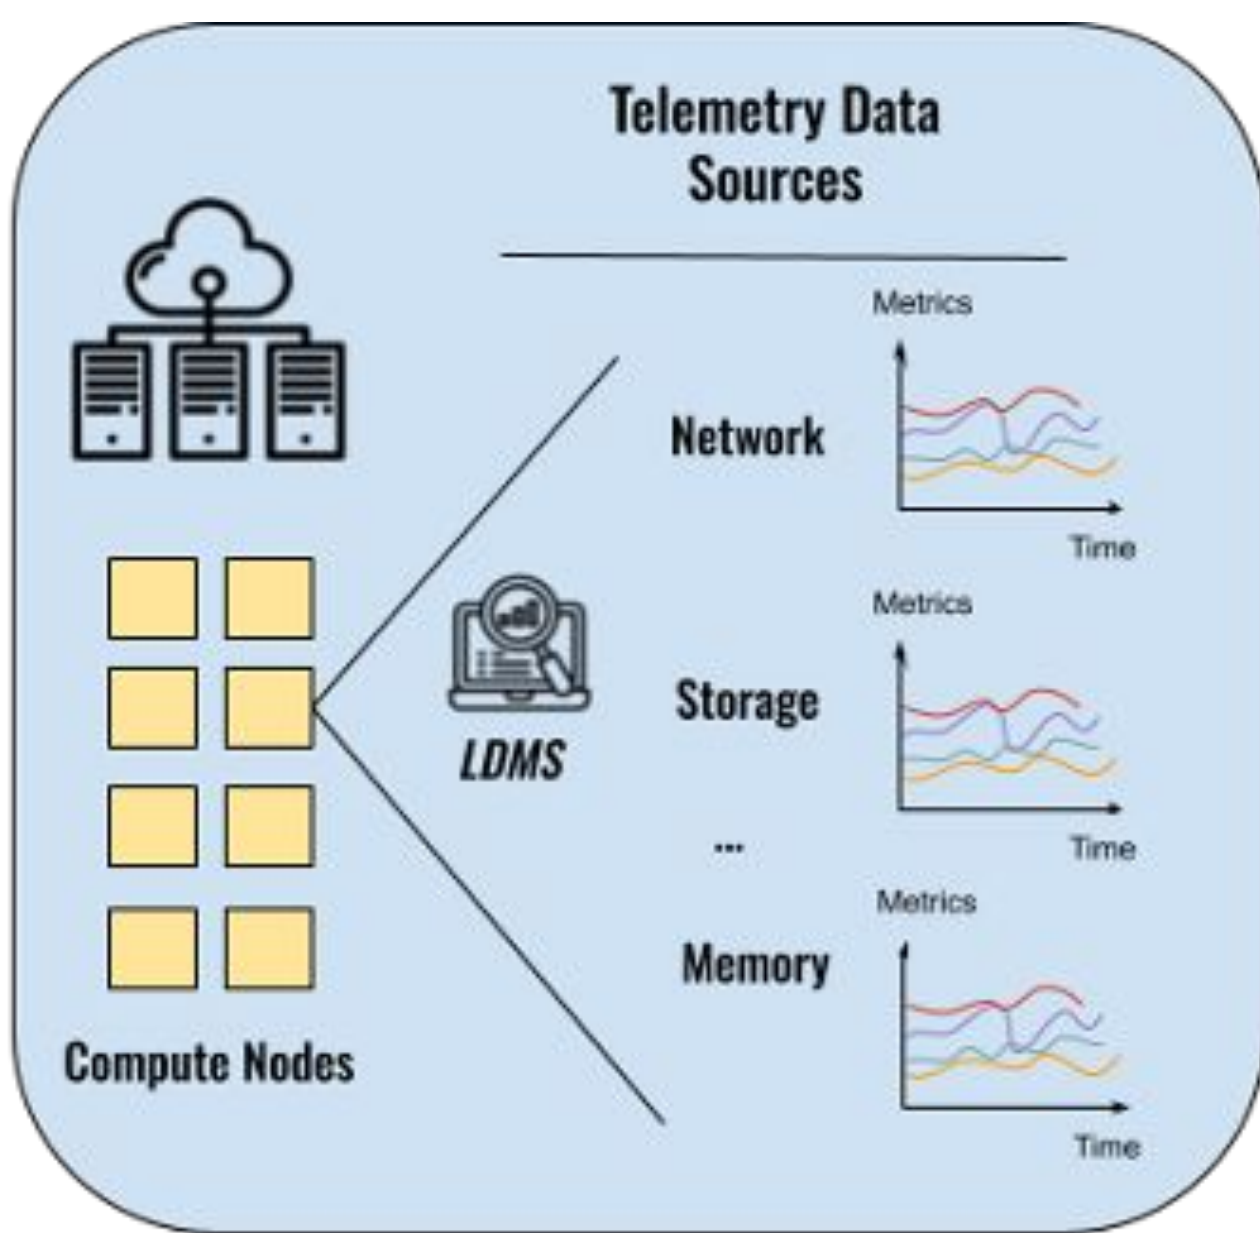
Background and Motivation

- Significant growth in the scale of large-scale computing systems
- Applications share network, memory, I/O, and computing resources
- This sharing leads to unpredictable application performance

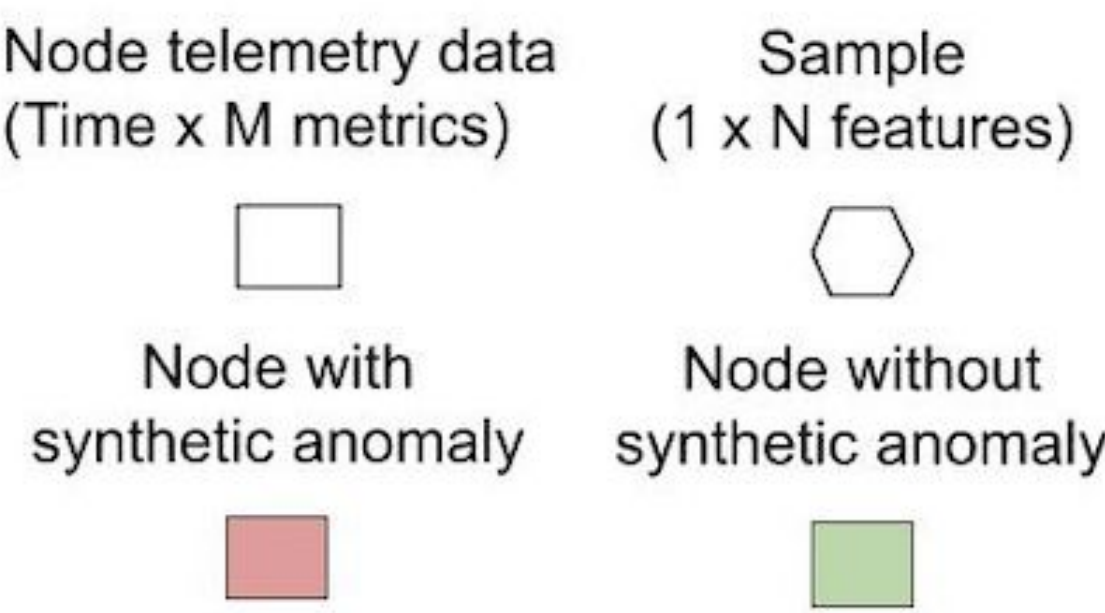


Problems caused by performance variations

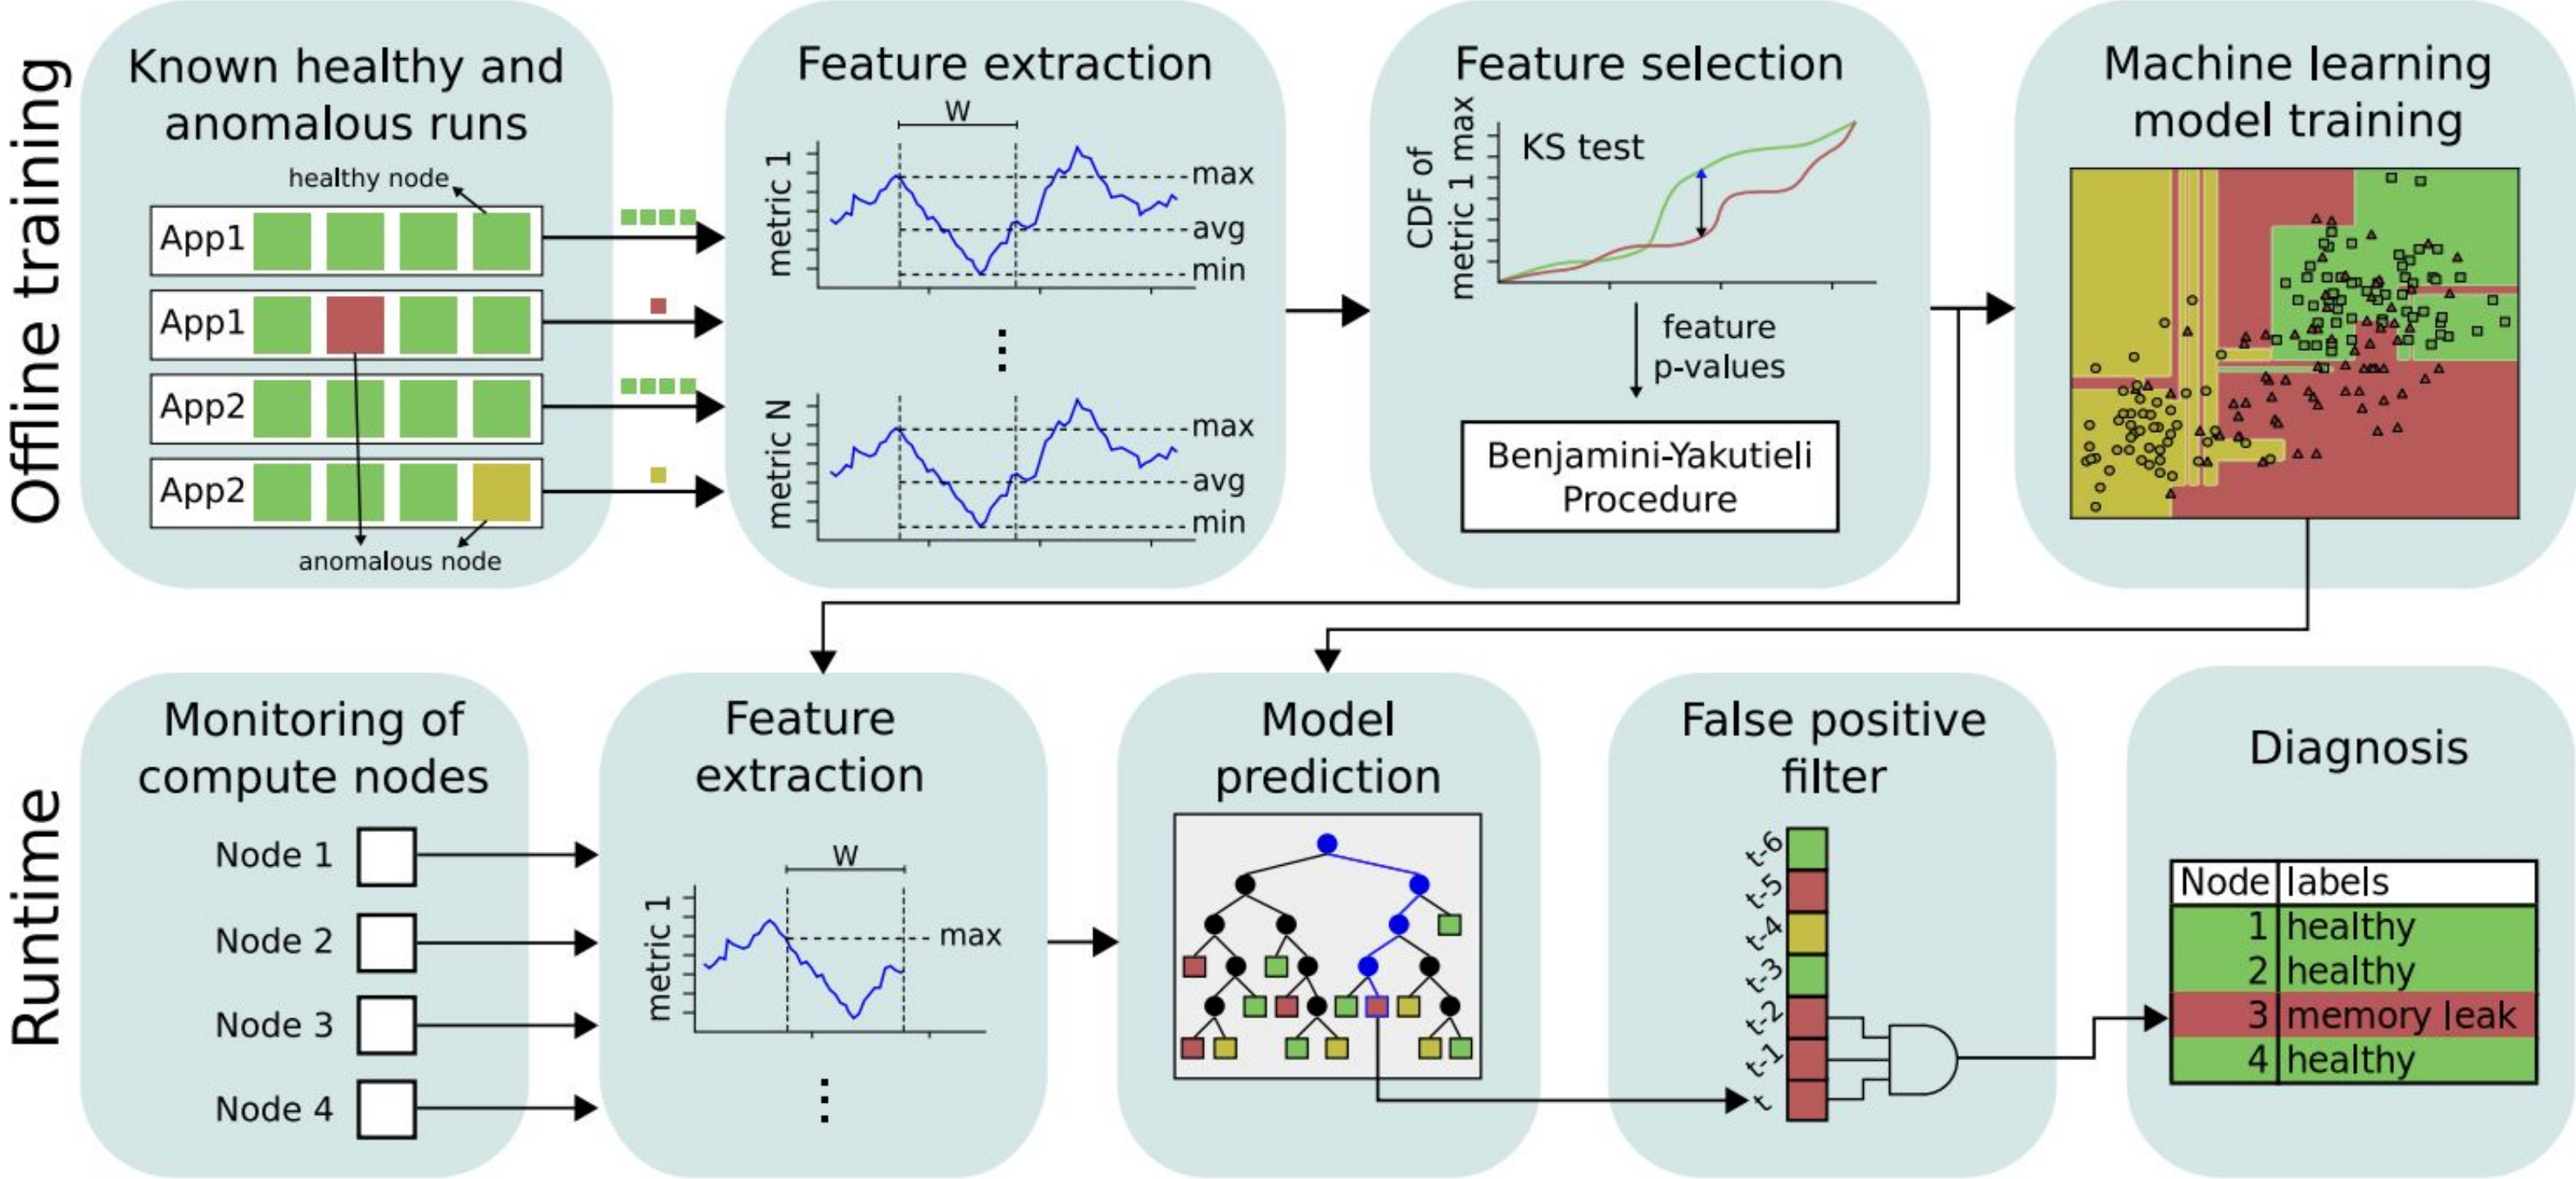
Telemetry Data Collection



Glossary:

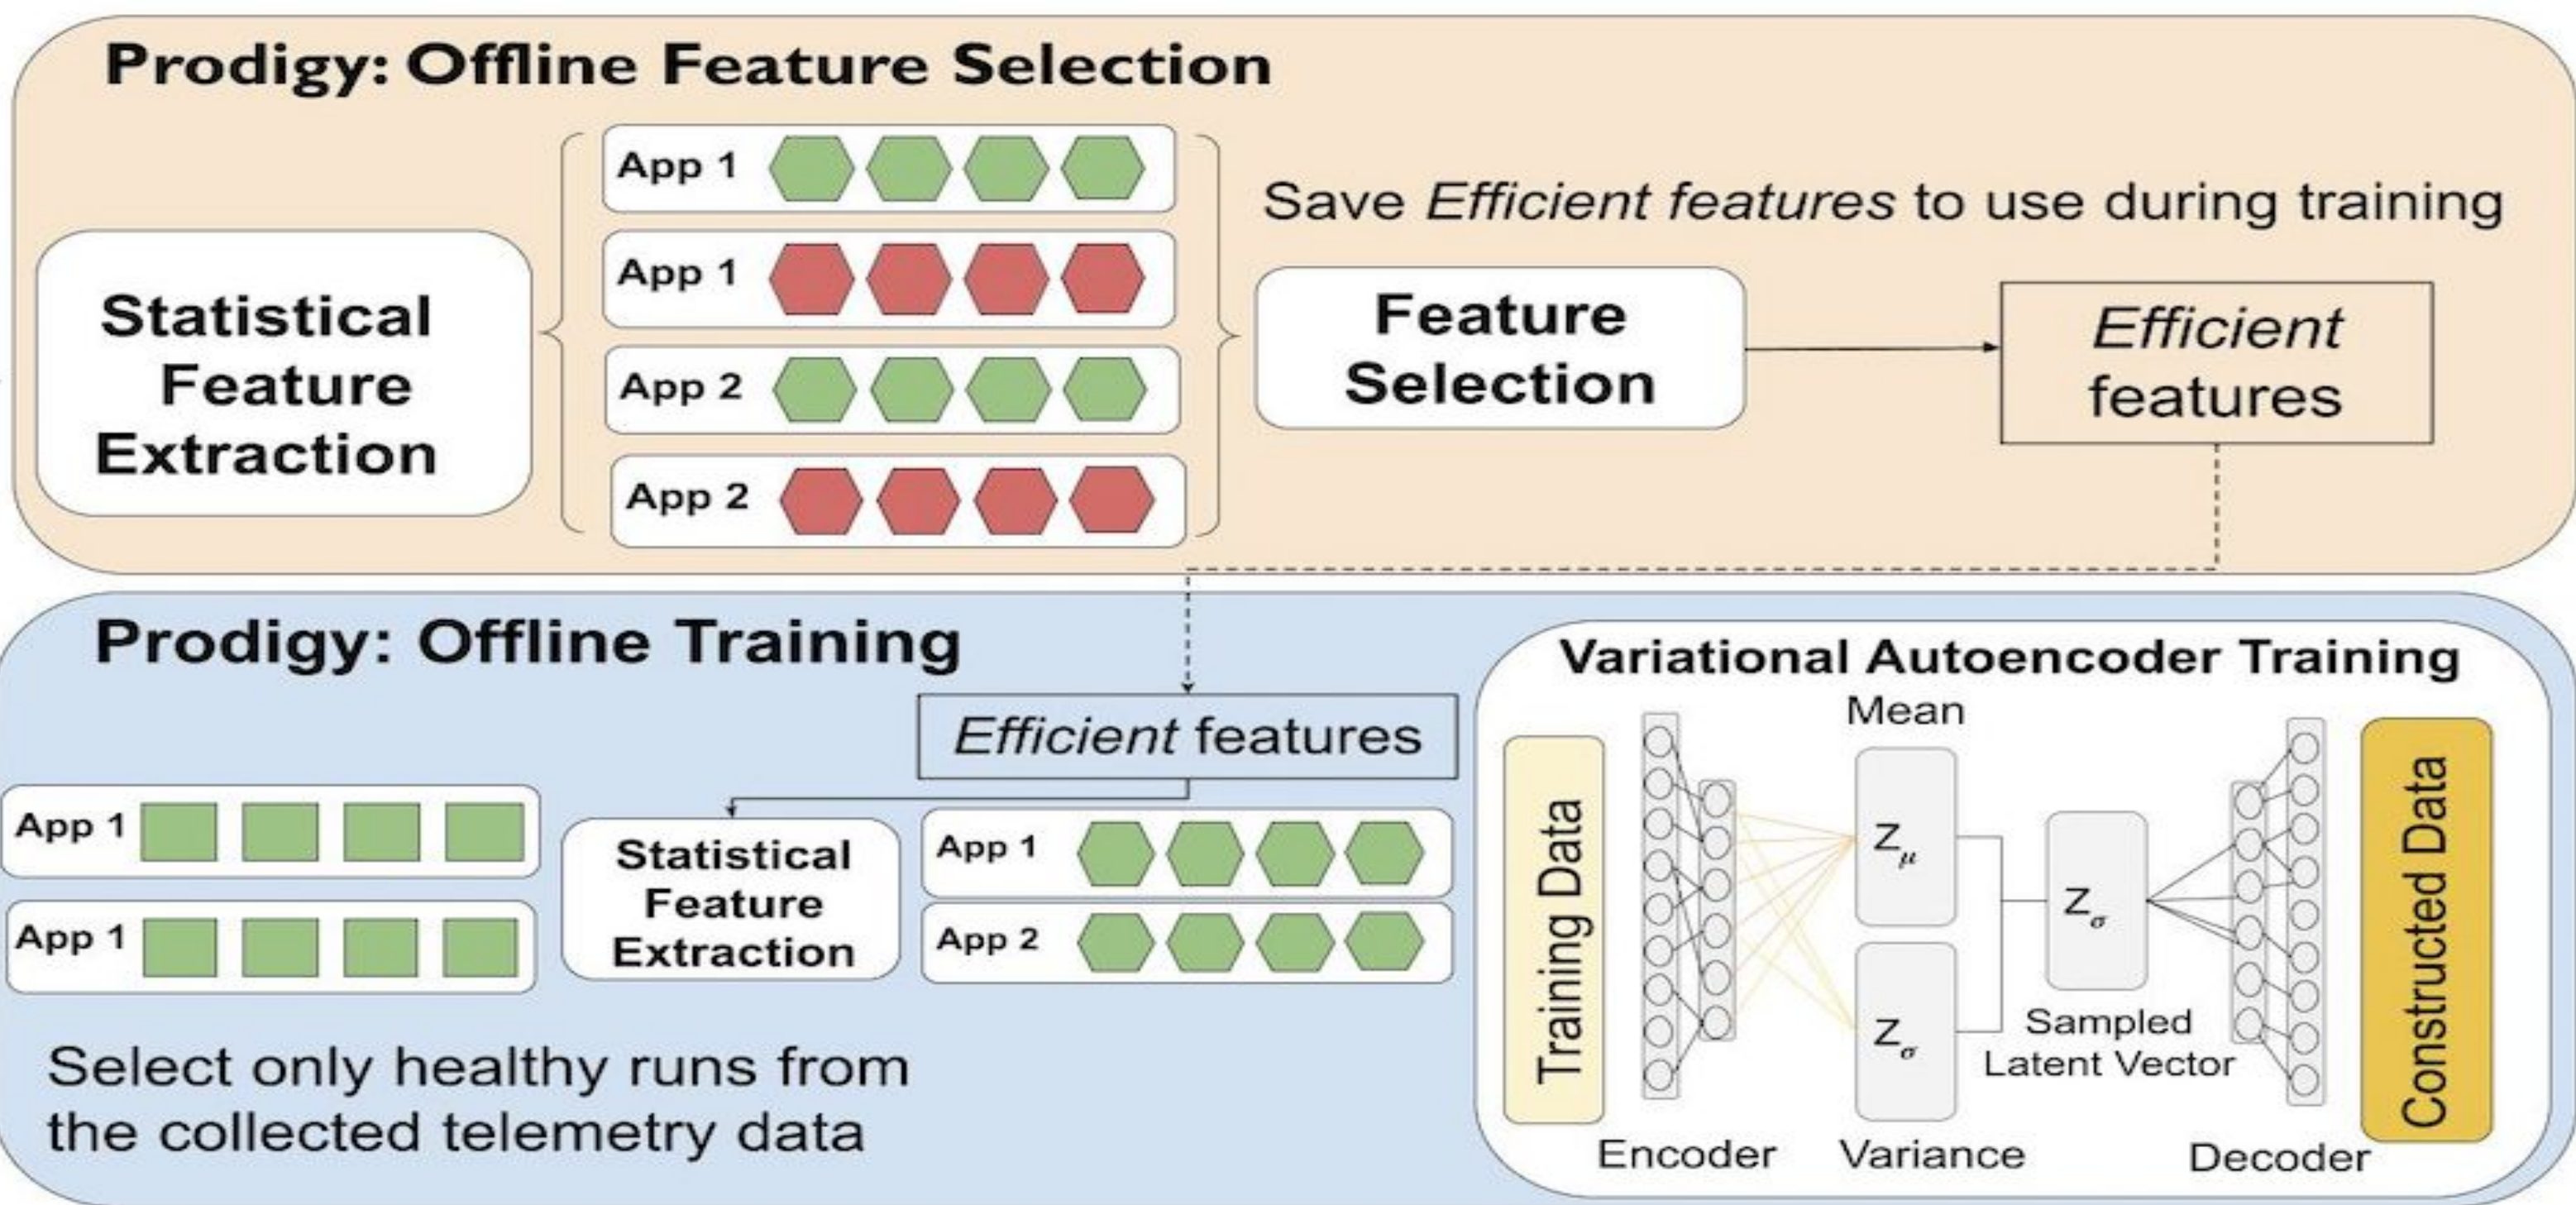


Supervised Framework Architecture



Overall pipeline for supervised anomaly diagnosis framework [2]

Unsupervised Framework Architecture

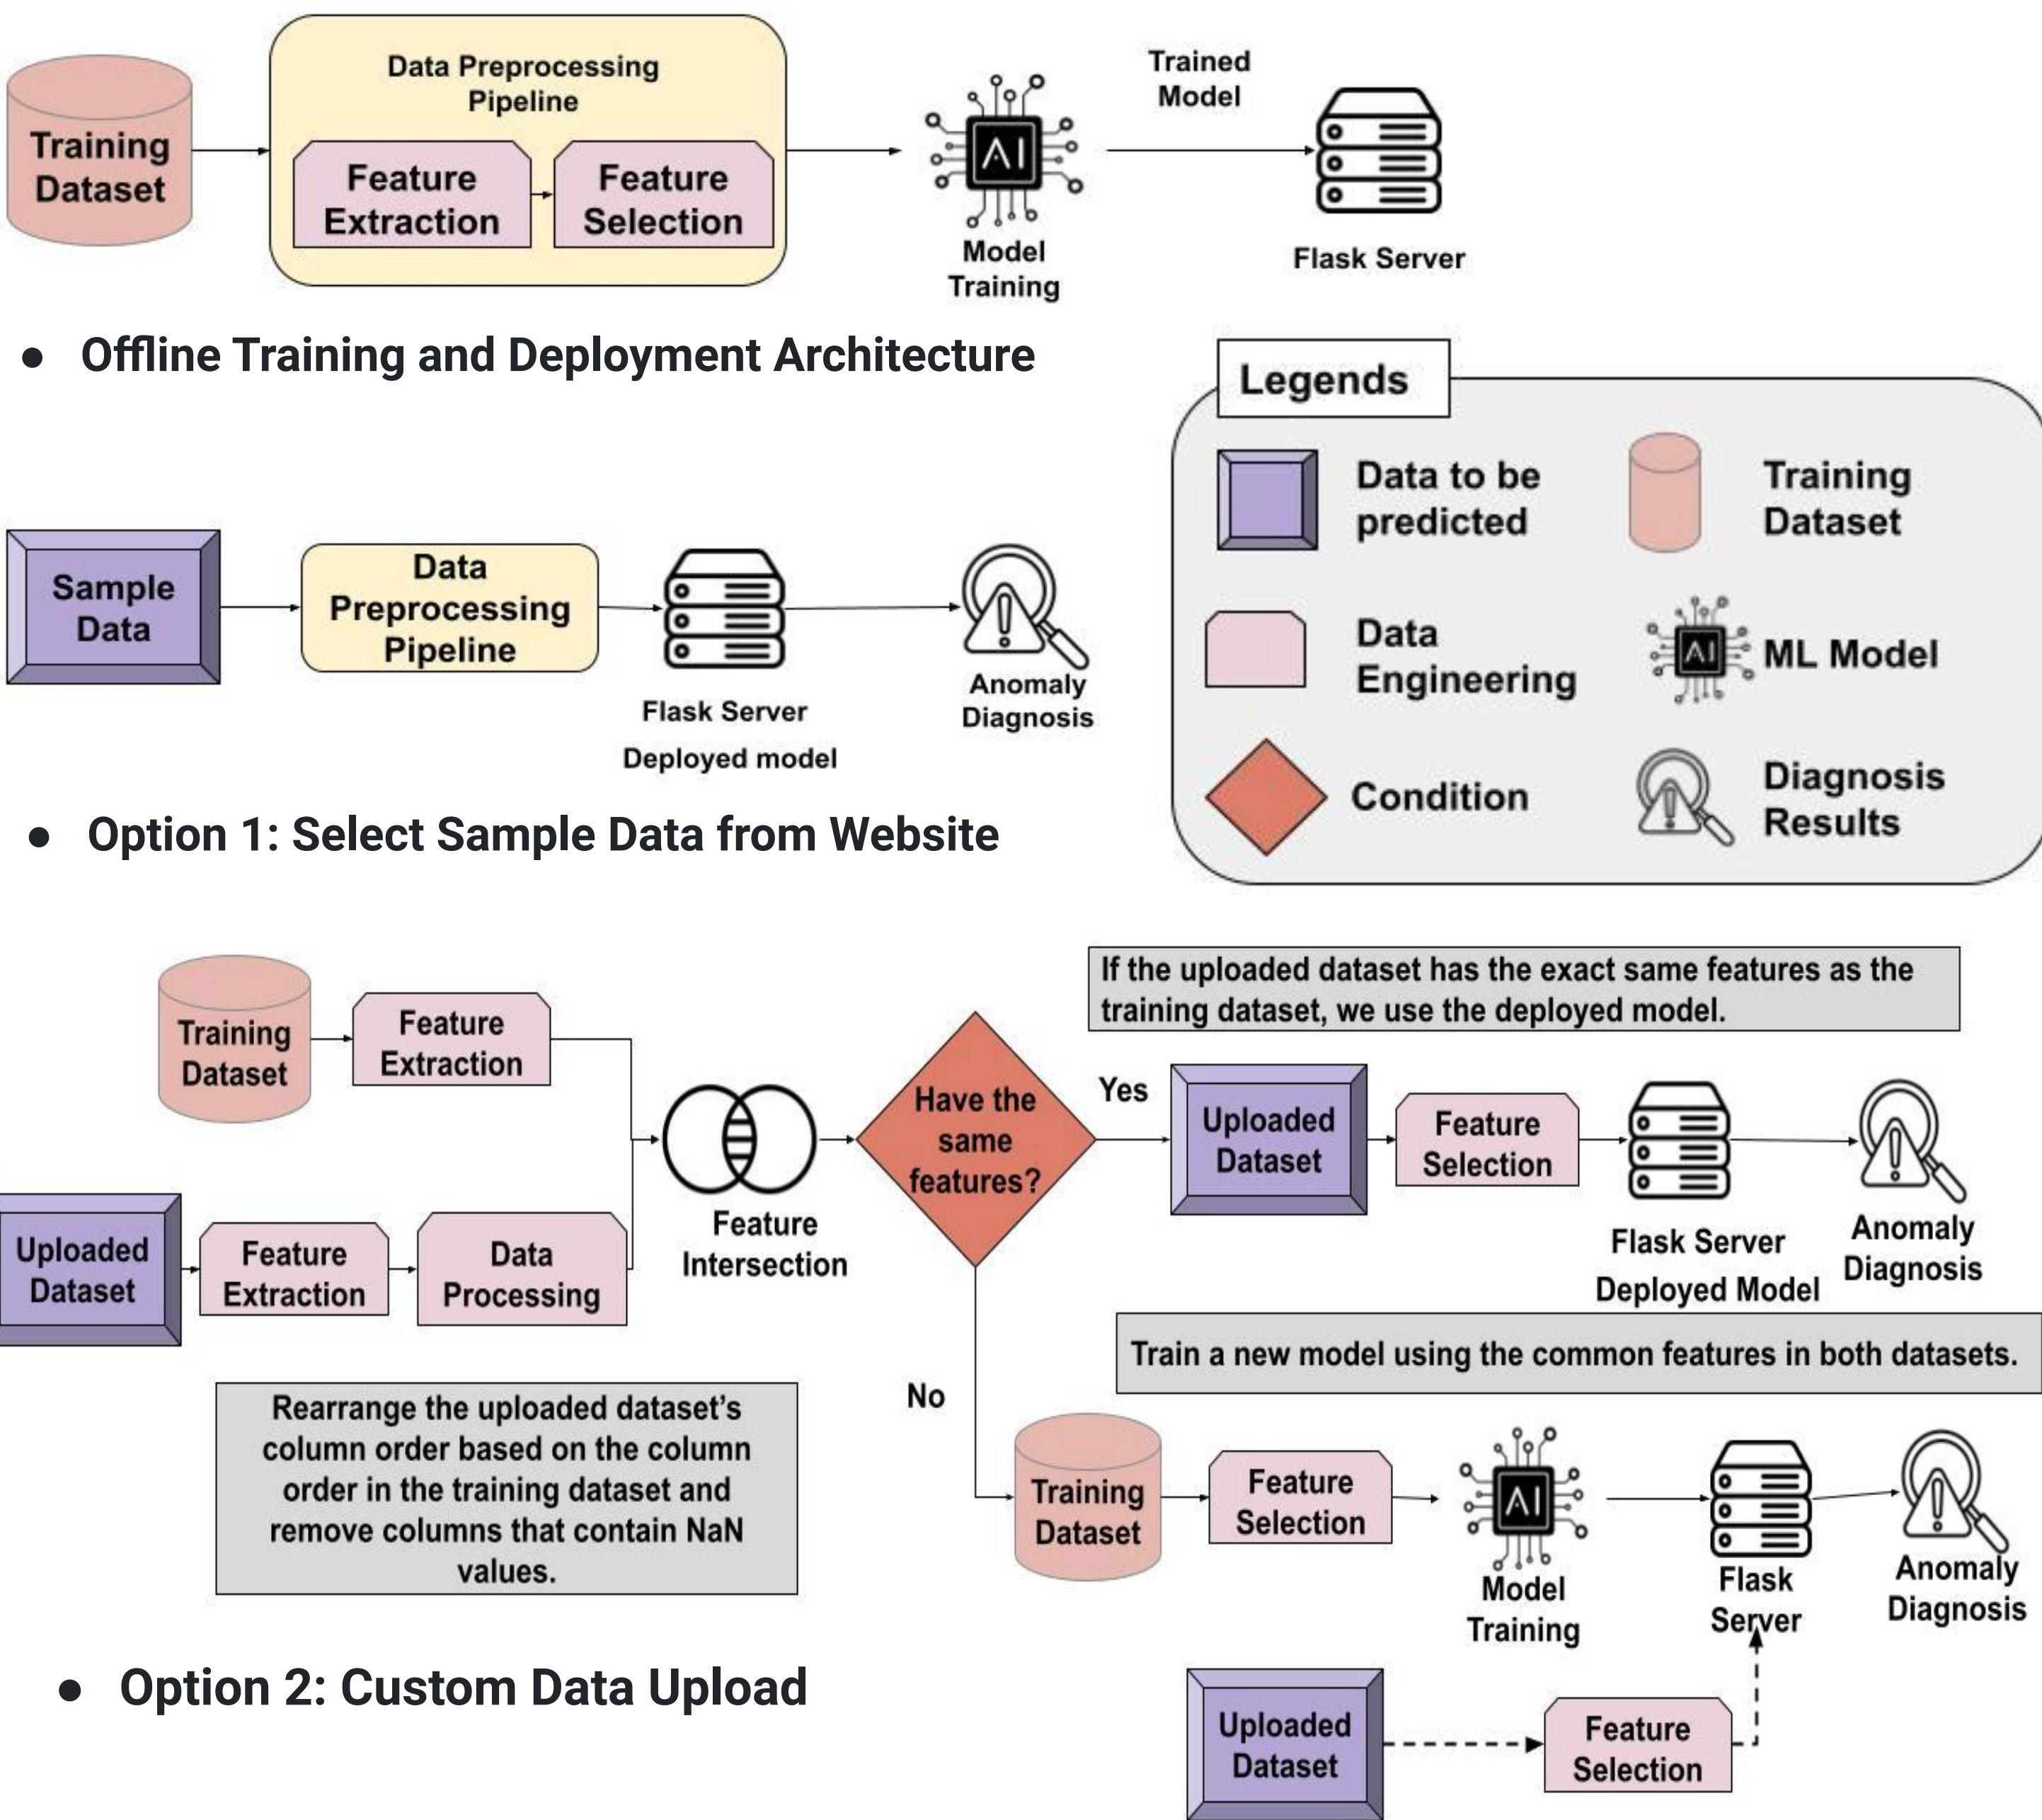


Prodigy: Unsupervised anomaly detection framework architecture [1]

Web-based Anomaly Diagnosis Framework

- Framework Access: AI4HPC offers both supervised and unsupervised anomaly detection frameworks in HPC systems.
- User Interaction: Enables training on custom datasets or application of pre-built models for analysis.
- Easy Setup: Docker integration for straightforward platform access and usage.

AI4HPC Architecture



Use Cases of AI4HPC

- Automated Anomaly Diagnosis: Utilizes AI for automated performance analysis, allowing users to upload and analyze their own dataset.
- Accessible ML Tools: Easy web access to ML models for analyzing HPC system data.
- User-Driven Enhancement: AI4HPC supports a feedback loop for users to contribute to the ongoing improvement of ML diagnostics.

References

- [1] B. Aksar, E. Sencan, B. Schwaller, O. Aaziz, V. J. Leung, J. Brandt, B. Kulis, M. Egele, and A. K. Coskun. Prodigy: Towards Unsupervised Anomaly Detection in Production HPC Systems. The International Conference on High Performance Computing, Network, Storage, and Analysis (SC 2023), Nov. 2023
- [2] O. Tuncer, E. Ates, Y. Zhang, A. Turk, J. Brandt, V. J. Leung, M. Egele, and A. K. Coskun. Online Diagnosis of Performance Variation in HPC Systems Using Machine Learning, in IEEE Transactions on Parallel and Distributed Systems (TPDS), vol. 30, no. 4, pp. 883-896, April 2019



Scan here for AI4HPC website

