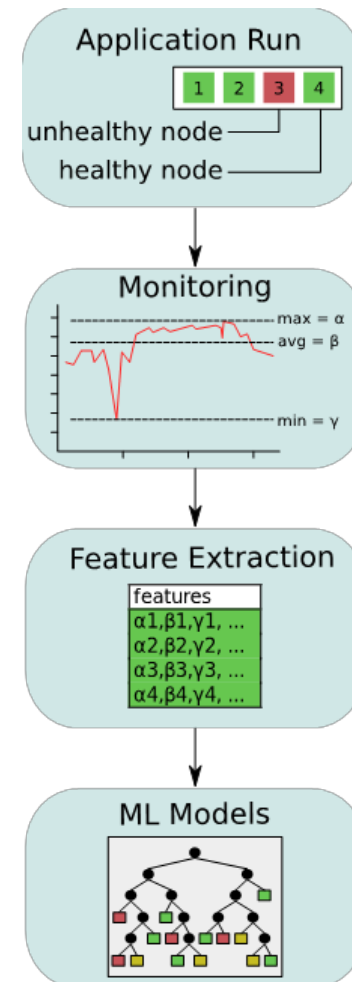
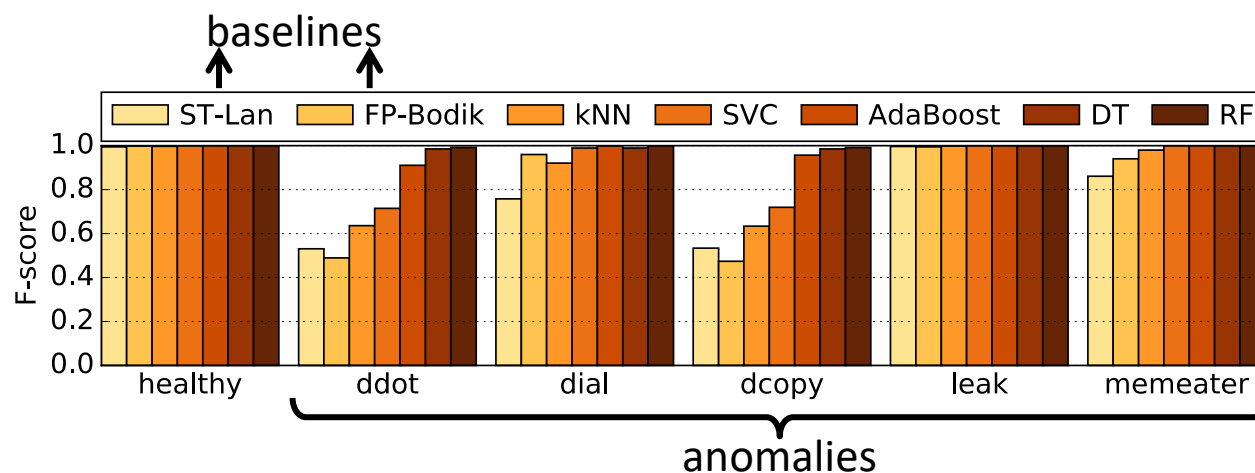


Detection and Diagnosis of Performance Variations*

- We proposed a novel framework to detect and diagnose anomalies (e.g., orphan processes, memory leaks, interference)
 - Based on resource usage data: 500 metrics collected every second
 - Uses concise statistical features
 - Has low storage and computation overhead
 - Supervised learning for classification
- Our framework outperforms existing anomaly detection methods
 - Evaluated on a Cray cluster and also on the MA Open Cloud
 - High diagnosis accuracy achieved, especially with Random Forest (RF)



* O. Tuncer, E. Ates, Y. Zhang, A. Turk, J. Brandt, V. J. Leung, M. Egele, A. K. Coskun. *Diagnosing Performance Variations in HPC Applications Using Machine Learning*. To appear in ISC-HPC'17. Gauss award recipient.