Thesis Defense

Computer Science Master's Program

"AIOps-Driven Adaptive Anomaly Detection in Evolving Cloud Environments Using Transfer Learning"

By **Mayur Shivakumar**

Abstract:

As cloud-based microservice architectures have become the foundation of contemporary enterprise solutions, performance interference, wherein co-located services compete for shared resources, remains a significant challenge. This phenomenon, often referred to as the noisy neighbor problem, manifests when one workload unexpectedly increases the CPU, memory, disk I/O, or network consumption, resulting in latency spikes or throughput degradation for other services. While existing isolation mechanisms (e.g., cgroups and QoS policies) provide some mitigation, they rarely prevent contention entirely, particularly in dynamic, rapidly evolving environments with frequent code deployments. This thesis proposes an AIOps-driven adaptive anomaly detection framework that integrates drift detection and Transfer Learning for real-time monitoring of containerized workloads. By focusing on operational metrics, including CPU, memory, I/O, and network usage, rather than solely application-level features, the system detects early performance interference signals. A key innovation is the application of the Kolmogorov-Smirnov (KS) test to identify statistically significant shifts in these resource distributions and automatically diOerentiate between benign updates and potentially disruptive anomalies. When the drift threshold is exceeded, selective retraining is initiated, ensuring that the models maintain accuracy without incurring the overhead of indiscriminate or excessively frequent retraining. Experimental validation employs Acme Air and the DeathStarBench Social Network, two diverse microservice platforms that exhibit varied resource consumption patterns. By systematically introducing new functionalities and monitoring workload evolution, the framework demonstrates how KS-based drift detection surpasses conventional staticthreshold methods in identifying early-stage noisy neighbor scenarios. Transfer Learning preserves prior knowledge while adapting rapidly to novel resource-usage profiles, oOering a cost-eOective approach for continuous, high-precision anomaly detection. In conclusion, this thesis bridges data drift detection and system-level performance monitoring within an MLOps context. It presents a scalable and proactive strategy for mitigating performance interference in multi-tenant, cloud-native environments, ultimately enhancing reliability and preserving service-level objectives.

Date: Friday, April 4th, 2025
Time: 3:30 PM – 5:30 PM
Location: Zoom
Committee: Dr. Mukherjee, Dr. Mukhopadhyay, and Dr. Kurfess