Title: Behavioral Side Effects of Model Compression on Structured Data

Target Students: B.Sc. and M.Sc. Students

Description:

Modern machine learning applications increasingly require deploying models on edge devices with strict computational and memory constraints. While compression techniques such as pruning and post-training quantization effectively reduce model size and inference latency, their impact extends far beyond simple accuracy trade-offs. This project investigates the behavioral side effects of model compression on time-series and tabular domains, where model reliability and safety are paramount. Depending on the student's degree level and focus, sub-projects can explore how weight sparsity and lower bit-precision degrade uncertainty calibration, shift anomaly detection thresholds, compromise feature attribution/explainability, or induce subgroup unfairness. While the primary emphasis is on temporal and tabular models, extensions to vision or language architectures are also possible for interested students.

Proposed Tasks:

The student will begin by selecting a specific behavioral axis (e.g., uncertainty calibration, anomaly thresholding, or feature attribution fidelity) and training baseline uncompressed models on representative benchmark datasets. They will then apply standard pruning (e.g., magnitude-based) and quantization protocols (e.g., INT8/INT4 PTQ or QAT) to evaluate how behavioral metrics degrade across increasing compression ratios. B.Sc. students will focus on empirical benchmarking and quantitative analysis across standard baselines. M.Sc. students will be expected to analyze the underlying mathematical failure modes (e.g., boundary deformation) and explore potential mitigation strategies, such as uncertainty-aware or calibration-aware pruning metrics.