Title: Algorithmic Recourse Under Model Compression and Low-Fidelity Proxies

Target Students: M.Sc. Students

Description:

Algorithmic recourse provides actionable recommendations for individuals seeking to reverse adverse automated decisions (e.g., loan denials or medical risk flags). However, real-world deployment pipelines often create a disconnect between the primary high-precision model hosted on central servers and compressed, lower-fidelity proxies deployed on client devices or made available for public auditing. This project explores the intersection of model compression (pruning and quantization) and algorithmic recourse optimization. The core objective is twofold: evaluating how compression-induced decision boundary shifts alter recourse costs or invalidate recommendations, and empirically exploring "low-fidelity algorithmic recourse." In this setting, we investigate how well recourse actions generated on a compressed, noisy proxy model actually transfer to and achieve approval on the central, uncompressed primary model.

Proposed Tasks:

The student will implement an experimental pipeline using standard tabular recourse benchmark datasets alongside established counterfactual search algorithms on Neural Networks. They will systematically prune and quantize target decision models to map how boundary shifts affect recourse cost inflation, search efficiency, and transferability. The student will then test empirical strategies to improve this transferability, such as introducing heuristic margin penalties on the compressed proxy to increase the likelihood of success on the true model. Formulating strict mathematical guarantees for this proxy-to-model transferability represents an open, highly advanced challenge, which serves as an optional next step for exceptionally interested students.