2025 – 01

Optimization of a Bandit-Guided Adaptive Sampling Framework

Yuan-chin Ivan Chang

May 22, 2025

Optimization of a Bandit-Guided Adaptive Sampling Framework

Yuan-chin Ivan Chang
Institute of Statistical Science, Academia Sinica, Taipei, Taiwan 11529

May 22, 2025

Abstract

Adaptive sampling in supervised learning is a crucial approach to optimize data efficiency, particularly in scenarios with limited labeling budgets or dynamic environments. This paper presents a novel framework that integrates multiple active sampling strategies—such as uncertainty, margin, and entropy sampling— with a multi-armed bandit (MAB) controller to dynamically select the most informative data acquisition method over time. The bandit selector supports various policies including Thompson Sampling, UCB, and Softmax exploration, enhanced with reward decay for non-stationary environments. We evaluate the system on a synthetic regression task with high-frequency nonlinearity, simulating 5000 trials to ensure statistical robustness. Results demonstrate consistent improvements in mean squared error (MSE) for dynamic strategy selection compared to fixed-policy baselines. Visualization of prediction confidence bands, strategy MSE trends, and arm selection frequencies confirms the adaptability and effectiveness of the proposed method. This modular design offers a general-purpose backbone for real-world active learning workflows under uncertainty and cost constraints.

Keywords and Phrases: Adaptive Sampling, Multi-Armed Bandits, Active Learning, Thompson Sampling, Exploration-Exploitation Tradeoff