ACLI: AN EFFICIENT CNN PRUNING FRAMEWORK EXPLOITING ADJACENT LAYER INTERDEPENDENCE AND Γ-WEAK SUBMODULARITY

Author's: Dhanush K, K. Sujith, N. Vanjulavalli

Abstract

Deep Convolutional Neural Networks (CNNs) have achieved remarkable performance across a wide range of computer vision tasks; however, their deployment in real-world edge and resource-constrained environments remains challenging due to high computational cost, memory footprint, and inference latency. To address these limitations, this paper proposes **ACLI (Adjacent Convolutional Layer Interdependence)**, a structured CNN pruning framework that integrates adjacent layer dependency modeling with a γ-weakly submodular optimization strategy. Unlike conventional pruning approaches that evaluate filter importance in isolation, ACLI explicitly models the interdependence between consecutive convolutional layers to better capture cross-layer feature propagation and redundancyThe proposed framework introduces a composite importance metric that combines filter saliencyestimated via gradient-based activation sensitivity—with an inter-layer influence score derived from the weight connectivity of subsequent convolutional layers. This dual-factor scoring mechanism ensures that pruning decisions consider not only local filter relevance but also their downstream contribution. To enhance selection robustness, ACLI formulates channel selection as a γ-weakly submodular maximization problem, enabling a greedy yet theoretically grounded pruning strategy that balances performance preservation and redundancy reduction. The γ parameter regulates the degree of diminishing returns during filter selection, promoting diversity among retained channels while preventing excessive information overlap.ACLI performs structured channel pruning, physically removing redundant filters and propagating dimensional consistency across subsequent layers. This results in a compact network architecture with reduced parameters and improved inference efficiency. A fine-tuning phase follows pruning to restore discriminative capacity and adapt to the restructured feature space. Experiments conducted on a synthetic multi-class image dataset demonstrate that ACLI achieves substantial parameter compression while maintaining competitive validation accuracy. Compared to baseline magnitude-based pruning, ACLI exhibits superior stability under aggressive pruning ratios due to its adjacency-aware scoring and diversity-aware selection mechanism.