

Provable Benefit of Cutout and CutMix for Feature Learning

Chulhee Yun

KAIST

Patch-level data augmentation techniques such as Cutout and CutMix have demonstrated significant efficacy in enhancing the performance of image-based tasks. However, a comprehensive theoretical understanding of these methods remains elusive. In this paper, we study two-layer neural networks trained using three distinct methods: vanilla training without augmentation, Cutout training, and CutMix training. Our analysis focuses on a feature-noise data model, which consists of several label-dependent features of varying rarity and label-independent noises of differing strengths. Our theorems demonstrate that Cutout training can learn low-frequency features that vanilla training cannot, while CutMix training can even learn rarer features that Cutout cannot capture. From this, we establish that CutMix yields the highest test accuracy among the three. Notably, our novel analysis reveals that CutMix training makes the network learn all features and noise vectors “evenly” regardless of the rarity and strength, which provides an interesting insight into understanding patch-level augmentation. This is joint work with Junsoo Oh (KAIST), to be presented at NeurIPS 2024 as a spotlight paper.