

Abstract

The go-to method applied to computer vision tasks involves convolutional neural networks. While computer vision has made great strides in utilizing convolutional neural networks there is still room for optimization. Many researchers still attempt to improve accuracy as well as improve computational requirements of the networks. ResNet is one such model which set out to improve the capabilities of these networks. In this work an extension of the residual-unit from ResNet is introduced, utilizing a self-attention mechanism. Using this mechanism, we have attempted to improve the classification accuracy of ResNet on the CIFAR-100 dataset, as well as attempting to use this mechanism as a tool to determine the most significant portions of a ResNet. Attempts to improve the performance of ResNet was inconclusive on CIFAR-100, but using this mechanism to determine effective regions of the network to prune has been successful.