

Deep Learning Action Recognition on Small Data: Transfer Learning the Stanford 40 Actions Dataset

Watson, Chandler

In the field of artificial intelligence, computers are exceptional at recognizing objects from images, but struggle with recognizing actions from images. In 2011, a human action dataset (Stanford 40 Actions) was released with fewer than 10,000 images. Much like an infant would struggle to recognize "walking" the first few times he/she observed it, computers have difficulty learning from small datasets. Previous research in action recognition on Stanford 40 Actions has had some success, but virtually none of it utilizes transfer learning (retraining an object recognizer into an action recognizer). This project, through a plethora of original scripts, applied transfer learning to successful object recognizers (convolutional neural networks) on Stanford 40 Actions and tested them both alone and in averaged clusters. Visualizations including low-dimensional embeddings and network weights plots were generated to inspect the object recognizers further. The most precise cluster of recognizers gave a mean average precision of 81.3%, whereas the largest value in the literature is 75.4% (a recent November 2015 paper). It was concluded that transfer learning is comparable to and outperforms previous algorithms in classifying Stanford 40 Actions.