



Super Resolution for Sports Images using a Generative Adversarial Network


James Liljenwall (jilijew@stanford.edu)

Motivation

Image technology keeps improving, and with it does the size of images. In part because of this, Super Resolution has become a major research field in Deep Learning. I wanted to test the performance of SRGAN, one of the most popular Super Resolution Algorithms, when trained on images just from sporting events, to try and optimize it for a genre of images

Models

SRGAN



C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, W. Shi, Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network, accepted at CVPR (oral), 2017.

Results

Hyper parameter

Batch Size	4	8	16	32
PSNR	50.04	50.308	50.687	50.25

From this, we can see that the best batch size was 16.

model	SRGAN (sport trained)	SRGAN (pre-trained)
PSNR	50.621	52.187

Data

My data came from images of sporting events, specifically, Rowing, Polo, Snow Boarding, Sailing, Badminton, Bocce, Croquet, and Rock Climbing. The images were then cropped and converted into low resolution versions



Error Analysis



Looking at the images, the images that are generated have a slightly sepia tone to them, which is likely a result of the training set being too small, with too much background. It is also slightly blurry, which is to be expected.

Conclusion/Future Work

Although the model performed worse than the pre-trained one, that was to be expected. It was still fairly close and the images were superior to the pixelated versions.

In the future, I would like to train on more data, and for longer, as I believe that including more images would allow the SRGAN to better fit to the task of sports images.