Video Superresolution

Student

October 2, 2017

Recent work has approached image superresolution using a variety of approaches [1, 2]. Other work has proposed novel loss functions for improved results [3]. How can we use these findings to do video superresolution? A simple baseline approach is to simply apply image superresolution independently to each frame. However, this is unlikely to result in a smooth video. Consecutive frames are inherently dependent on each other, and we would like to take advantage of this.

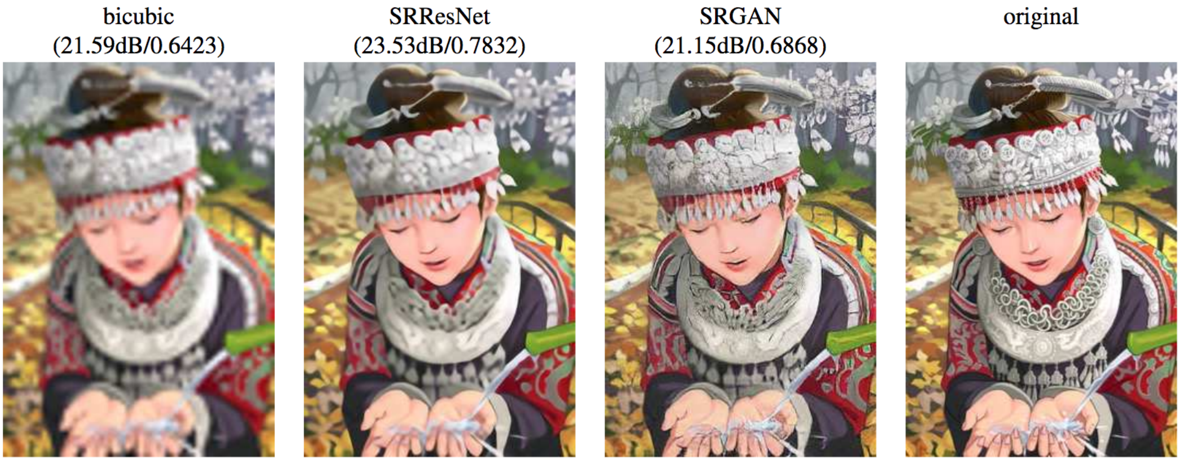

Figure 1: Figures from [2] comparing various recent methods for image superresolution.

In this project, you will investigate how we can intelligently apply image superresolution techniques to video. Rather than applying superresolution to each frame independently, you might want to add some sort of temporal smoothing across frames, or incorporate additional constraints into the loss function.

Fortunately, procuring high quality training data will not be difficult for this task. You may simply downsample videos of your choice to use for training. We would like this to be an exploratory project in which you propose an original model and/or loss function to improve upon the baseline.

References

[1] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 1874–1883.

[2] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, "Photo-realistic single image super-resolution using a generative adversarial network," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[3] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *European Conference on Computer Vision*, Springer, 2016, pp. 694–711.