

Photo Uncrop

Bad Photographer

October 2018

To help photographers worldwide shoot better photos, you will create a computer vision system that uncrops photos. Given a photo, the computer program will extrapolate the pixel values by some given amount. See the figure below.

Collecting training data for this project is easy. You can just take any image dataset and randomly crop the photos to create pairs of inputs/targets for training a model.

There are a couple of ways to do this. For example, you could learn to retrieve other photos that are similar and stitch them together [3]. You will probably need to experiment with different approaches to blend images together [2].

You could also experiment with using a generative model, such as a GAN, to automatically extrapolate images [1].



Figure 1: Can you create a computer vision system to uncrop a photo?

References

- [1] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2536–2544, 2016.
- [2] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. *ACM Transactions on graphics (TOG)*, 22(3):313–318, 2003.
- [3] Qi Shan, Brian Curless, Yasutaka Furukawa, Carlos Hernandez, and Steven M Seitz. Photo uncrop. In *European Conference on Computer Vision*, pages 16–31. Springer, 2014.