NVIDIA researchers raise the bar on image inpainting

April 26 2018, by Nancy Owano

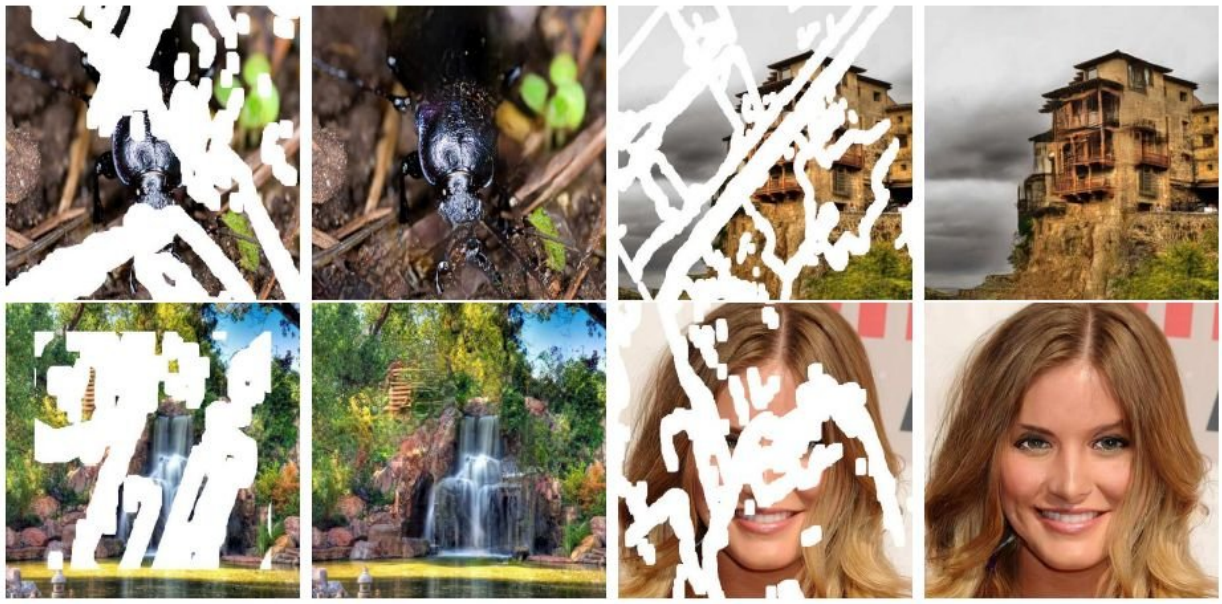

Masked images and corresponding inpainted results using our partialconvolution based network. Credit: arXiv:1804.07723 [cs.CV]

For those not yet familiar with photo reconstruction tools, the title of this paper on arXiv might be totally puzzling: "Image Inpainting for Irregular Holes Using Partial Convolutions." The research, though, from a NVIDIA team, points the way to exciting improvements in store for those who must perform image editing and who desire good results.

Image inpainting is all about filling in holes in an image. It can be used to take out image content that is not wanted, while filling in the space with plausible imagery. Back to the title of their paper, the team explored their take of an improved process, which could be implemented in photo editing software.

Over at the NVIDIA Developer News Center is an overview about their research. They came up with a method that one serve to (1) edit images or (2) reconstruct a corrupted image, one that has holes or is missing pixels. And when they said "edit," that includes removing content and filling in the holes.

The video clearly shows how interesting this can get, presenting sets of photo before and after the whiting-out process kicks in. The scenes for one is a rock outdoors and another is a library indoors. Yet another set shows humans' faces, including a woman, young males and an elderly man.

What's their work all about?

"Researchers from NVIDIA, led by Guilin Liu, introduced a state-of-the-art deep learning method that can edit images or reconstruct a corrupted image, one that has holes or is missing pixels. The method can also be used to edit images by removing content and filling in the resulting holes." That is according to the video notes.

At play were two phases, the training phase and the testing phase.

To prepare to train their neural network, the team first generated masks of random streaks and holes of arbitrary shapes and sizes for training, said the news center report. Categories were devised based on sizes relative to the input image, to improve reconstruction accuracy. The neural network training involved the generated masks to images from

ImageNet, Places2 and CelebA-HQ datasets.

"During the training phase, holes or missing parts are introduced into complete training images from the above datasets, to enable the network to learn to reconstruct the missing pixels. During the testing phase, different holes or missing parts, not applied during training, are introduced into the test images in the dataset, to perform unbiased validation of reconstruction accuracy."

Why their work stands out: "To the best of our knowledge, we are the first to demonstrate the efficacy of deep learning image inpainting models on irregularly shaped holes."

The researchers were aware of existing deep learning-based image inpainting methods. These were using "a standard convolutional network over the corrupted image, using convolutional filter responses conditioned on both valid pixels as well as the substitute values in the masked holes (typically the mean value)."They said this "often leads to artifacts such as color discrepancy and blurriness. Post-processing is usually used to reduce such artifacts, but are expensive and may fail."

They said they were proposing partial convolutions—by that "the convolution is masked and renormalized to be conditioned on only valid pixels."

They showed qualitative and quantitative comparisons with other methods to validate their approach, and they stated that their model "outperforms other methods for irregular masks."

The paper's authors are Guilin Liu, Fitsum Reda, Kevin Shih, Ting-Chun Wang, Andrew Tao and Bryan Catanzaro.

More information: Image Inpainting for Irregular Holes Using Partial

Convolutions, arXiv:1804.07723 [cs.CV] arxiv.org/abs/1804.07723

Abstract
Existing deep learning based image inpainting methods use a standard convolutional network over the corrupted image, using convolutional filter responses conditioned on both valid pixels as well as the substitute values in the masked holes (typically the mean value). This often leads to artifacts such as color discrepancy and blurriness. Post-processing is usually used to reduce such artifacts, but are expensive and may fail. We propose the use of partial convolutions, where the convolution is masked and renormalized to be conditioned on only valid pixels. We further include a mechanism to automatically generate an updated mask for the next layer as part of the forward pass. Our model outperforms other methods for irregular masks. We show qualitative and quantitative comparisons with other methods to validate our approach.

news.developer.nvidia.com/new- … h-realistic-results/

Citation: NVIDIA researchers raise the bar on image inpainting (2018, April 26) retrieved 23 February 2024 from https://techxplore.com/news/2018-04-nvidia-bar-image-inpainting.html