New algorithm repairs corrupted digital images in one step

December 5 2017

TOP: The research team artificially degraded a stock image, deliberately introducing blur, noise and other imperfections. BOTTOM: The research team's new image repair algorithm automatically returned the image to near-original quality. Credit: Don Cochran, Kodak Lossless True Color Image Suite

From phone camera snapshots to lifesaving medical scans, digital images play an important role in the way humans communicate information. But digital images are subject to a range of imperfections such as blurriness, grainy noise, missing pixels and color corruption.

A group led by a University of Maryland computer scientist has designed a new algorithm that incorporates artificial neural networks to simultaneously apply a wide range of fixes to corrupted digital images. Because the algorithm can be "trained" to recognize what an ideal, uncorrupted image should look like, it is able to address multiple flaws in a single image.

The research team, which included members from the University of Bern in Switzerland, tested their algorithm by taking high-quality, uncorrupted images, purposely introducing severe degradations, then using the algorithm to repair the damage. In many cases, the algorithm outperformed competitors' techniques, very nearly returning the images to their original state.

The researchers presented their findings on December 5, 2017, at the 31st Conference on Neural Information Processing Systems in Long Beach, California.

"Traditionally, there have been tools that address each problem with an image separately. Each of these uses intuitive assumptions of what a good image looks like, but these assumptions have to be hand-coded into

the algorithms," said Matthias Zwicker, the Reginald Allan Hahne Endowed E-Nnovate Professor in Computer Science at UMD and senior author of the research presentation. "Recently, artificial neural networks have been applied to address problems one by one. But our algorithm goes a step further—it can address a wide variety of problems at the same time."

Artificial neural networks are a type of [artificial intelligence algorithm](link) inspired by the structure of the human brain. They can assemble patterns of behavior based on input data, in a process that resembles the way a human brain learns new information. For example, human brains can learn a new language through repeated exposure to words and sentences in specific contexts.

Zwicker and his colleagues can "train" their algorithm by exposing it to a large database of high-quality, uncorrupted images widely used for research with artificial neural networks. Because the algorithm can take in a large amount of data and extrapolate the complex parameters that define images—including variations in texture, color, light, shadows and edges—it is able to predict what an ideal, uncorrupted image should look like. Then, it can recognize and fix deviations from these ideal parameters in a new image.

"This is the key element. The algorithm needs to be able to recognize a good image without degradations. But for an image that is already degraded, we can't know what this would look like," said Zwicker, who also has an appointment at the University of Maryland Institute for Advanced Computer Studies (UMIACS). "So instead, we first train the algorithm on a database of high-quality images. Then we can give it any image and the algorithm will modify the imperfections."

Zwicker noted that several other research groups are working along the same lines and have designed algorithms that achieve similar results.

Many of the research groups noticed that if their algorithms were tasked with only removing noise (or graininess) from an image, the algorithm would automatically address many of the other imperfections as well. But Zwicker's group proposed a new theoretical explanation for this effect that leads to a very simple and effective algorithm.

"When you have a noisy image, it is randomly shifted or jittered away from a high-quality image in all possible dimensions. Other degradations, such as blurring for example, diverge from the ideal only in a subset of dimensions," Zwicker explained. "Our work revealed how fixing noise will bring all dimensions back in line, allowing us to address several types of other degradations, like blurring, at the same time."

Zwicker also said that the new algorithm, while powerful, still has room for improvement. Currently, the algorithm works well for fixing easily recognizable "low-level" structures in images, such as sharp edges. The researchers hope to push the algorithm to recognize and repair "high-level" features, including complex textures such as hair and water.

"To recognize high-level features, the algorithm needs context to understand what is in the image. For example, if there is a face in an image, it's likely that the pixels near the top are probably hair," Zwicker said. "It's like assembling a jigsaw puzzle. If you're only looking at one piece, it's hard to place that part of the image in context. But once you find where the piece belongs, it's much easier to recognize what the pixels represent. It's quite clear that this approach can be pushed much further still."

The research paper, "Deep Mean-Shift Priors for Image Restoration," Siavash A. Bigdeli, Meiguang Jin, Paolo Favaro and Matthias Zwicker, was presented on December 5, 2017 at the 31st Conference on Neural Information Processing Systems in Long Beach, California.

More information: Deep Mean-Shift Priors for Image Restoration, arXiv:1709.03749 [cs.CV] arxiv.org/abs/1709.03749