Team accelerates rendering with AI

June 30 2017

Modern films and TV shows are filled with spectacular, computer-generated sequences which are computed by rendering systems that simulate the flow of light in a 3D scene. However, computing many light rays is an immensely labor-intensive and time-consuming process. The alternative is to render the images using only a few light rays, but this shortcut results in inaccuracies that show up as objectionable noise in the final image.

Researchers from Disney Research, Pixar Animation Studios, and the University of California, Santa Barbara have developed a new technology based on artificial intelligence (AI) and deep learning that eliminates this noise and thereby enables production-quality rendering at much faster speeds.

Specifically, the team used millions of examples from the Pixar film Finding Dory to train a deep learning model known as a Convolutional Neural Network. Through this process, the system learned to transform the noisy images into noise-free images that resemble those computed with significantly more light rays. Once trained, the system was successfully able to remove the noise on test images from entirely different films, such as Pixar's latest release, "Cars 3," and their upcoming feature "Coco," even though they had completely different styles and color palettes

"Noise is a really big problem for production rendering," said Tony DeRose, head of research at Pixar. "This new technology allows us to automatically remove the noise while preserving the detail in our

scenes."

The work presents a significant step forward over previous, state-of-the-art denoising methods which often left artifacts or residual noise that required artists to either render more light rays or to tweak the denoising filter to improve the quality of a specific image. Disney and Pixar plan to incorporate the technology in their production pipelines to accelerate the movie-making process.

"Other approaches for removing image noise have grown increasingly complex, with diminishing returns," said Markus Gross, vice president for research at Disney Research. "By leveraging deep learning, this work presents an important step forward for removing undesirable artifacts from animated films."

The work will be presented in July at the ACM SIGGRAPH 2017 conference, the premier venue for technical research in computer graphics. To facilitate further exploration of this exciting area, the team will make their code and trained weights available to the research community.

Provided by Disney Research