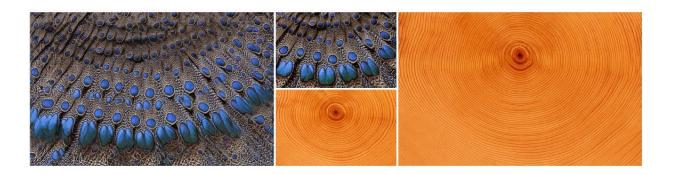# AI-based framework creates realistic textures in the virtual world

July 16 2018

Sample images of 'texture synthesis' using a unique artificial intelligence-based technique that trains a network to learn to expand small textures into larger ones. This data-driven method leverages an AI technique called generative adversarial networks (GANs) to train computers to expand textures from a sample patch into larger instances that best resemble the original sample. Credit: Zhen Zhu, Xiang Bai, Dani Lischinski, Daniel Cohen-Or, and Hui Huang

Many designers for the virtual world find it challenging to design efficiently believable complex textures or patterns on a large scale. Indeed, so-called "texture synthesis," the design of accurate textures such as water ripples in a river, concrete walls, or patterns of leaves, remains a difficult task for artists. A plethora of non-stationary textures in the "real world" could be re-created in gaming or virtual worlds, but the existing techniques are tedious and time-consuming.

To address this challenge, a global team of computer scientists has developed a unique artificial intelligence-based technique that trains a network to learn to expand small textures into larger ones. The researchers' data-driven method leverages an AI technique called generative adversarial networks (GANs) to train computers to expand textures from a sample patch into larger instances that best resemble the original sample.

"Our approach successfully deals with non-stationary textures without any high level or semantic description of the large-scale structure," says Yang Zhou, lead author of the work and an assistant professor at Shenzhen University and Huazhong University of Science & Technology. "It can cope with very challenging textures, which, to our knowledge, no other existing method can handle. The results are realistic designs produced in high-resolution, efficiently, and at a much larger scale."

The basic goal of example-based texture synthesis is to generate a texture, usually larger in size than the input, that closely captures the visual characteristics of the sample input—yet not entirely identical to it—and maintains a realistic appearance. Examples of non-stationary textures include textures with large-scale irregular structures, or ones that exhibit spatial variance in certain attributes such as color, local orientation, and local scale. In the paper, the researchers tested their method on such complex examples as peacock feathers and tree trunk ripples, which are seemingly endless in their repetitive patterns.

Zhou and his collaborators, including Zhen Zhu and Xiang Bai (Huazhong University), Dani Lischinski (The Hebrew University of Jerusalem), Daniel Cohen-Or (Shenzhen University and Tel Aviv University), and Hui Huang (Shenzhen University), will present their work at SIGGRAPH 2018, held 12-16 August in Vancouver, British Columbia. This annual gathering showcases the world's leading

professionals, academics, and creative minds at the forefront of computer graphics and interactive techniques.

Their method involves training a generative network, called generator, to learn to expand (i.e., double the spatial extent of) an arbitrary texture block cropped from an exemplar, so that the expanded result is visually similar to a containing exemplar block of the appropriate size (two times larger). The visual similarity between the automatically expanded block and the actual containing block is assessed using a discriminative network (discriminator). As typical of GANs, the discriminator is trained in parallel to the generator to distinguish between actual large blocks from the exemplar and those produced by the generator.

Says Zhou, "Amazingly, we found that by using such a conceptually simple, self-supervised adversarial training strategy, the trained network works near-perfectly on a wide range of textures, including both stationary and highly non-stationary textures."

The tool is meant to assist texture artists in video game design, virtual reality, and animation. Once the self-supervised adversarial training takes place for each given texture sample, the framework may be used to automatically generate extended textures, up to double the original sample size. Down the road, the researchers hope their system will be able to actually extract high-level information of textures in an unsupervised fashion.

Additionally, in future work, the team intends to train a "universal" model on a large-scale texture dataset, as well as increase user control. For [texture](#) artists, controlled synthesis with user interaction will likely be even more useful since artists tend to manipulate the textures for their own design.

For the full paper and video, visit the team's project page.