

Researchers use Web images to add realism to edited photos

July 10 2007

Computer graphics researchers at Carnegie Mellon University have developed systems for editing or altering photographs using segments of the millions of images available on the Web.

Whether adding people or objects to a photo, or filling holes in an edited photo, the systems automatically find images that match the context of the original photo so they blend realistically. Unlike traditional photo editing, these results can be achieved rapidly by users with minimal skills.

“We are able to leverage the huge amounts of visual information available on the Internet to find images that make the best fit,” said Alexei A. Efros, assistant professor of computer science and robotics. “It’s not applicable for all photo editing, such as when an image of a specific object or person is added to a photo. But it’s good enough in many cases,” he added. “Why Photoshop if you can ‘photoswap’ instead?”

Efros and his colleagues will present papers on two related systems at the Association for Computing Machinery’s Special Interest Group on Graphics and Interactive Techniques (SIGGRAPH) annual conference Aug. 5–9 in San Diego.

One system, called Photo Clip Art (graphics.cs.cmu.edu/projects/photoclipart/), was developed with graduate students Jean-François Lalonde and Derek Hoiem, and with

Carsten Rother, John Winn and Antonio Criminisi of Microsoft Research Cambridge. It uses thousands of labeled images from a Web site called LabelMe as clip art that can be added to photos. A photo showing a vacant street, for instance, might be populated with images of people, vehicles and even parking meters derived from the LabelMe database (labelme.csail.mit.edu/).

To make the resulting image appear as realistic as possible, the system analyzes the original photo to estimate the camera angle and lighting conditions, and then looks in the clip art library for an object — a car, for instance — that matches those criteria. The user need only identify the horizon in the original photo to orient the system. Using previously developed Carnegie Mellon technology for analyzing the geometric context of a photo, the system can then place the object within the scene, adjusting its size as necessary to put it in proportion to other objects of equal distance from the camera.

“Matching an object with the original photo and placing that object within the 3-D landscape of the photo is a complex problem,” said Lalonde, who led development of the system. “But with our approach, and a lot of clip art data, we can hide the complexity from the user and make the process simple and intuitive.”

The other system, called Scene Completion (graphics.cs.cmu.edu/projects/scene-completion/), was developed by graduate student James Hays, another member of Efros’ research team. It draws upon millions of photos from the Flickr Web site to fill in holes in photos. Some of the holes might be from damage to a physical photograph, but more often they are created when an editor cuts out part of an image to eliminate an unsightly truck from a picturesque street scene, or removing a passerby from a group shot of friends. Photo editors often try to fill in those holes with sections derived elsewhere in the same image, but Efros said that a better match can often be found in

a different photo.

The system looks for image segments that match the colors and textures that surround the hole on the original photo. It also looks for image segments that make sense contextually — in other words, it wouldn't put an elephant in a suburban backyard or a boat in a desert.

In the case of well-photographed cities or popular tourist attractions, Efros said, the system might get lucky and find a photo of the same scene on the Web. In other cases, it might offer a number of possible images that could fill in the hole. A retaining wall edited out of one photo, for instance, might be replaced by the image of a building, a grassy slope or a rock outcropping. The system typically gives the user 20 different choices for filling in the hole.

The success of this approach depends on the number of photos available to the system, Hays said. "We saw a dramatic improvement when we moved from a database of 10,000 images to two million images," he noted. "And that is just a tiny fraction of the hundreds of millions of images already available on sites like Picasa and Flickr. We have tons of photos from which to choose."

Source: Carnegie Mellon University

Citation: Researchers use Web images to add realism to edited photos (2007, July 10) retrieved 2 May 2024 from <https://phys.org/news/2007-07-web-images-realism-photos.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.
