TJHSST Computer Systems Lab Senior Reasearch Project
Dynamic Image Resizing

Patrick Elliott

June 10, 2008

1 Abstract

Effective resizing of images should not only use geometric constraints, but consider the image content as well. This project uses a simple image operator called seam carving that supports content-aware image resizing for both reduction and expansion. A seam is an optimal 8-connected path of pixels on a single image from top to bottom, or left to right, where optimality is defined by an image energy function. By repeatedly carving out or inserting seams in one direction we can change the aspect ratio of an image. By applying these operators in both directions, we can retarget the image to a new size. This method was explained in detail in the paper "Content-Aware Image Resizing."

The selection and order of seams protect the content of the image, as defined by the energy function. Seam carving can also be used for image content enhancement and object removal. We support various visual saliency measures for defining the energy of an im- age, and can also include user input to guide the process. By storing the order of seams in an image we create multi-size images, that are able to continuously change in real time to fit a given size.

2 Introduction

The diversity and versatility of display devices today imposes new demands on digital media. For instance, designers must create different alternatives for web-content and design different layouts for different devices. Moreover, HTML, as well as other standards, can support dynamic changes of page layout and text. Nevertheless, up to date, images, although being one of the key elements in digital media, typically remain rigid in size and cannot deform to fit different layouts automatically. Other cases in which the size, or aspect ratio of an image must change, are to fit into different displays such as cell phones or PDAs, or to print on a given paper size or resolution.

Standard image scaling is not sufficient since it is oblivious to the image content and typically can be applied only uniformly. Cropping is limited since it can only remove pixels from the image edges. More effective resizing can only be achieved by considering the image content and not only geometric constraints.

I propose a simple image operator, seam-carving, that can change the size of an image by gracefully carving-out or inserting pixels in different parts of the image. Seam carving uses an energy function defining the importance of pixels. A seam is a connected path of low energy pixels crossing the image from top to bottom, or from left to right. By successively removing or inserting seams we can reduce, as well as enlarge, the size of an image in both directions. For image reduction, seam selection ensures that while preserving the image structure, we remove more of the low energy pixels and fewer of the high energy ones. For image enlarging, the order of seam insertion ensures a balance between the original image content and the artificially inserted pixels. These operators produce, in effect, a content-aware resizing of images.

We illustrate the application of seam carving and insertion for aspect ratio change. Furthermore, by storing the order of seam removal and insertion operations, and carefully interleaving seams in both vertical and horizontal directions we define multi-size images. Such images can continuously change their size in a content-aware manner. A designer can author a multi-size image once, and the client application, depending on the size needed, can resize the image in real time to fit the exact layout or the display.

3 Background

Edge detection is being researched heavily in modern times. Many teams of researchers are trying to allow computers to see and identify objects. But there is also much research being conducted concerning editing modifying images. There is one project called PhotoSynth that is trying to take a large amount of images from the web, and from them, create a 3D model of whatever the images are of. There is also another project that is very similar to what I am attempting to do, although I have some ideas for my project that they have not yet implemented.

4 Development

I will be using C for all of my programming. My approach to content-aware resizing is to remove pixels in a judicious manner. Therefor, the question is how to chose the pixels to be removed? Intuitively, my goal is to remove unnoticeable pixels that blend with their surroundings. This leads to the following simple energy function: Each pixel's value is the difference in intensity with the pixel's neighbors.

Given an energy function, assume we need to reduce the image width. One can think of several strategies to achieve this. For instance, an optimal strategy to preserve energy (i.e., keep pixels with high energy value) would be to remove the pixels with lowest energy in ascending order. This destroys the rectangular shape of the image, because we may remove a different number of pixels from each row. If we want to prevent the image from breaking we can remove an equal number of low energy pixels from every row. This preserves the rectangular shape of the image but destroys the image content by creating a zigzag effect. To preserve both the shape and the visual coherence of the image we can use auto-cropping. That is, look for a sub-window, the size of the target image, that contains the highest energy. However, if two important objects are on either side of the target image, this wont achieve the desired effect.

Another possible strategy somewhat between removing pixels and cropping is to remove whole columns with the lowest energy. Still, artifacts might appear in the resulting image. Therefore, we need a resizing operator that will be less restrictive than cropping or column removal, but can preserve the image content better than single pixel removals. This leads to my strategy of seam carving.

Rather than removing scattered pixels or entire columns, my program removes seams (vertical paths of pixels within one column of the pixel above it) of the lowest intensity. The optimal seam can be found using dynamic programming. The first step is to traverse the image from the second row to the last row and compute the cumulative minimum energy M for all possible connected seams for each entry (i, j): M(i, j) = e(i, j) + min(M(i-1, j-1),M(i-1, j),M(i-1, j+1)) where the function e is the intensity of the given pixel. At the end of this process, the minimum value of the last row in M will indicate the end of the minimal connected vertical seam. Hence, in the second step we backtrack from this minimum entry on M to find the path of the optimal seam. The definition of M for horizontal seams is similar.

To enlarge an image we insert new artificial seams to the image. Hence, to enlarge the size of an image I by one we compute the optimal vertical (horizontal) seam s on I and duplicate the pixels of s by averaging them with their left and right neighbors (top and bottom in the horizontal case).

5 Results

The program is capable of shrinking the image in both the horizontal and vertical directions and enlarging in the horizontal direction. With the newest method of finding the optimal pixels to remove (the cumulative sums of the gradient image) the final images are usually very realistic. It is difficult to notice that it has been modified at first glance, which is the goal of the project.

6 Limitations

All the tests in this project were computed automatically, but my method does not work automatically on all images. This can be corrected by adding higher level cues, either manual or automatic. For example, sometimes images with faces become distorted, but combined with a face detector it would get much better results.

Other times, not even high level information can solve the problem. We can characterize two major factors that limit our seam carving approach. The first is the amount of content in an image. If the image is too condensed, in the sense that it does not contain ”less important” areas, then any type of content-aware resizing strategy will not succeed. The second type of limitation is the layout of the image content. In certain types of images, albeit not being condensed, the content is laid out in a manner that prevents the seams from bypassing important parts.