Seam Carving for Content-Aware Image Resizing

Matthew Thorne

Shai Avidan and Ariel Shamir. Seam carving for content-aware image resizing. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, page 10, New York, NY, USA, 2007. ACM

The paper "Seam Carving for Content-Aware Image Resizing" by Avidan and Shamir proposes a method of resizing images such that the image content is not distorted as would be the case by a simple image scaling operation. To achieve this, they introduce the concept of a *seam* - a low energy chain of pixels stretching from one side of an image to the other than can be removed (or inserted) with minimal visual impact. This is useful, for example, when retargetting an image to different display devices.

The authors first consider the case of shrinking an image by removing low energy pixels from an image (those whose removal would have least visual impact). They use the energy function: $e_1(\mathbf{I}) = |\frac{\partial}{\partial x}\mathbf{I}| + |\frac{\partial}{\partial y}\mathbf{I}|$ but also observe that other functions can be used. Since naively removing low energy pixels from an image will cause artefacts, the authors present their *seam carving* operator which removes pixels in a way that seeks to preserve image content.

If $\mathbf{I}$ is an *n x m* image, then a *vertical seam* is defined as: $\mathbf{s^x} = \{s_i^x\}_{i=1}^n = \{(x(i), i)\}_{i=1}^n$ such that $\forall i|x(i) - x(i-1)| \leq 1$ and *x* is a mapping $x : [1, \ldots, n] \rightarrow [1, \ldots, m]$. Informally, this means a vertical seam is an 8-connected path of pixels from the top to the bottom of the image containing exactly one pixel per row. A horizontal seam can be similarly defined. To reduce the width (or height) of an image, the vertical (or horizontal) seam with least energy is removed. This seam is found by minimizing $s^* = \min_s E(\mathbf{s}) = \min_s \sum_{i=1}^n e(\mathbf{I}(s_i))$ where *e* is an energy function. This minimization problem is solved using dynamic programming.

Seam removal can be repeated iteratively to reduce the width (or height) of the image by a specific number of pixels. Simultaneously reducing the width and height requires seams in both directions to be removed, but what ordering of horizontal and vertical seam removal gives the best results? If the target image size is $n'xm'$ with $n' < n$ and $m' < m$ this can be decided by optimizing: $\min_{s^x, s^y, \alpha} \sum_{i=1}^k E(\alpha_\mathbf{i} \mathbf{s_i^x} + (1-\alpha_i)\mathbf{s_i^y})$ where $k = r+c, r = (m-m')$ and $c = (n - n')$ using α_i as a parameter that deterines which type of seam was removed at step *i*. $\alpha_i \in \{0, 1\}$, $\sum_{i=1}^k \alpha_i = r$, $\sum_{i=1}^k (1 - \alpha_i) = c$. This optimization is also computed using dynamic programming.

Seam carving can also be used to enlarge images. This is done by inserting 'artificial' seams into an image by finding the optimal seam in one direction and duplicating the pixels along it by averaging them with their neighbours. Since repeating this process will tend to insert the same seam over and over, the authors identify the first k seams that would be removed and duplicate them to arrive at the final enlarged image.

To speed up seam carving for large images, the authors present a multi-size representation for images in which, for each pixel, the index of the seam that would have removed that pixel is stored for both the horizontal and vertical directions. To reduce the width of an image by w pixels, those with seam index w or less can quickly be removed. Care must be taken when supporting resizing in both directions, however, since horizontal and vertical seams can conflict with each other. The authors observe that an easy way to solve this is to use seams in one direction and simply remove rows or columns in the other direction. They also describe how to apply seams in both directions by applying a consistency constraint.

The authors also present two other applications for their seam carving technique: using it to amplify the content in an image (such as enlarging a certain region) and to remove objects from images. While interesting, these are more of an aside to the main idea of the paper and aren't given many details.

The main idea - resizing an image using seams of low energy pixels - is interesting and well presented. though it does have some flaws, as the authors point out, since it doesn't work automatically for all images - faces in particular seem to cause problems, though using a face detector can solve these. It would have been interesting to see more failure cases and whether or not something like an edge detector or image segmentation could be used to assist with or detect problem areas. They also very briefly mention the possibility of using manual interaction to correct problems but offer no suggestions of how this might be approached. These could be possibilities for future work.

In the paper, the authors also briefly touch on using other energy functions, but do not provide much details. They observe that no one energy function works well on all images, but they have similar "range" for resizing. I would have liked to see a little more discussion on this matter as it may have been informative to know if there were certain classes of images that tended to be more suited for certain energy functions than others.