hw1
Brian Burns
CS376 Computer Vision, Spring 2011
Dr. Kristen Grauman

4. [10 points] Make some change to the way the energy function is computed (i.e., filter used, its parameters, or incorporating some other a prior knowledge). Display the result and explain the impact on the results for some example.

Tried using a non-smoothing edge filter on an image (ie [-1 0 1] instead of the Sobel filter), but there didn't seem to be any change in the energy map - the bright areas just looked a little dimmer. So...

The painting Christina's World seems to have an area of pink (the woman's dress) which would be good to avoid. So I tried making a mask for the pink hue, and adding that to the energy function. However, a lot of the image actually has the same hue, just darker intensity, eg in the lower right corner, so as the image got reduced down further, the woman got chopped up, when the grass would have been a better target.

Here's the original energy map -

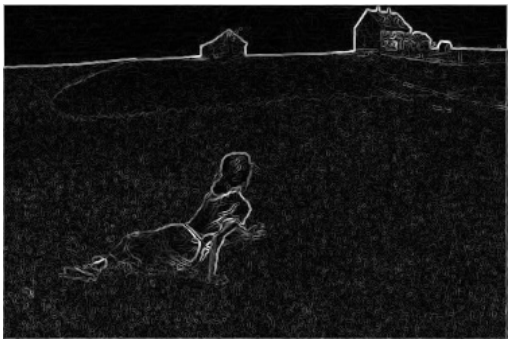

Then the hue mask - the hue of her dress is 3-25 out of 255, so made a simple filter based on those hues. Note the areas on the edges that had the same hue.

Then just added this to the original image map (with each scaled to have a maximum value of 1) -

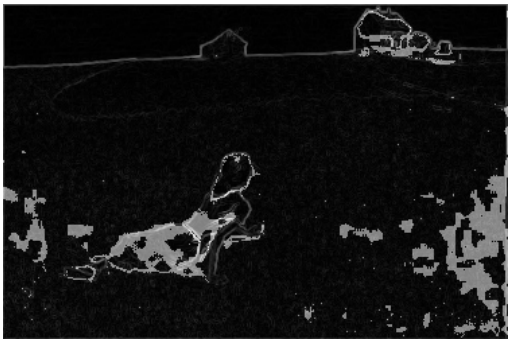

Then tried resizing the image - it did seem to work better on removing 100 columns. The result of using the new energy function is on the right - notice that her figure is less distorted -

But removing another 100 columns winds up eating into her figure, more so than the original energy function does -

So adding the hue mask helped for a while, then the red hues on the edges started edging out the girl.

A better approach would probably be to use the full rgb value of the pink in the dress, and define some mask based on closeness to that value, instead of just the hue value.

5. [30 points] Now, for the real results! Use your system with different kinds of images and seam combinations, and see what kind of interesting results it can produce. The goal is to form some perceptually pleasing outputs where the resizing better preserves content than a blind resizing would, as well as some examples where the output looks unrealistic or has artifacts.

i. prague

(a) original image

(b) seam carved image

(c) resized using imresize

(d) input size: 640x480, output size: 540x480
(e) removed 100 vertical seams

(f) Note the slight distortion in the central building in the seam-resized version - it would be better to go through the trees on the left, but the energy function would assign them too much weight, due to their texture. A better approach might try to account for common textures like trees or grass by giving them lower energy levels.

(a) original image

(b) seam carved image

(c) resized using imresize

(d) input size: 775x769, output size: 675x769
(e) removed 100 vertical seams
(f) Note the distortion of the main tree in the seam-carved version - the tree has been almost eliminated, because it is an area of low contrast.

iii. appleworks

(a) original image

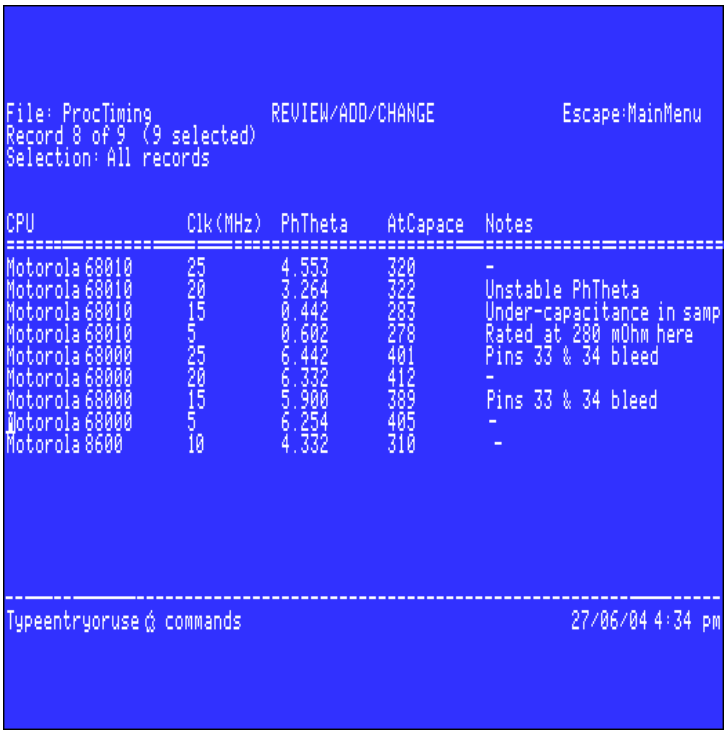

(b) seam carved image

(c) resized using imresize

(d) input size: 722x526, output size: 522x526
(e) removed 200 vertical seams
(f) This is an area which the seam-carving algorithm is helpful - instead of squeezing the text until it's almost unreadable, as the normal resize method does, the seam-carving algorithm preferentially removes blank areas, leaving the text more legible.
(g) Image is a screenshot from Appleworks, an old Apple II application.

(a) original image

(b) seam carved image

(c) resized using imresize

(d) input size: 520x250, output size: 370x250
(e) removed 150 vertical seams
(f) This was chosen to be a bad example, as there's no alternative for the algorithm but to carve into the people, and our eyes are very sensitive to distortions in people's figures and faces. But it almost works - perhaps with a skin hue sensitive energy function, the faces could be better preserved - the result might look a bit cartoonish, with normal heads on stick figure bodies, but their faces would at least be preserved.
(g) Image is of the cast of Lost, from ABC Television.

v. austin

(a) original image

(b) seam carved image

(c) resized using imresize

(d) input size: 1024x341, output size: 512x341
(e) removed 512 vertical seams
(f) This seam carving worked pretty well, though I was expecting more to get taken out of the buildings. But the grass in the foreground adds a lot to the energy function, so the buildings are protected. This is another case where recognizing a texture like grass would be helpful.
(g) Image is by Joergen Geerds, from http://newyorkpanorama.com/tag/austin/

vi. nyc

(a) original image

(b) seam carved image

(c) resized using imresize

(d) input size: 1024x324, output size: 512x324

(e) removed 512 vertical seams
(f) And here's a case where the normal resizing method works MUCH better. The seam-carving takes out different areas of the bridges, resulting in severe distortions of their linearity. This is a case where combining seam-carving with normal resizing might be helpful - eg remove some of the bland areas like the central buildings, and bridge towers, but when a certain cumulative energy threshold is reached, switch to normal resizing operations.
(g) Image is by Joergen Geerds, from http://newyorkpanorama.com

III. [OPTIONAL] Extra credit [up to 10 points each, max possible 20 points extra credit]

I just made a movie showing the seam carving for an image, with 100 columns being removed. The program **resize.m** created and saved all the images, then **makemovie.m** created an avi file from them. The avi file was uncompressed though, at 78MB, so used Windows Movie Maker to convert it to a wmv file, which is 1MB.

jack.wmv

or

http://www.youtube.com/watch?v=W0WhcbF6xTw