


CS 378 Computer Vision
Problem set 1
Out: Tuesday Sept 8
Due: Monday Sept 21, by 11:59 PM

See the end of this document for submission instructions.

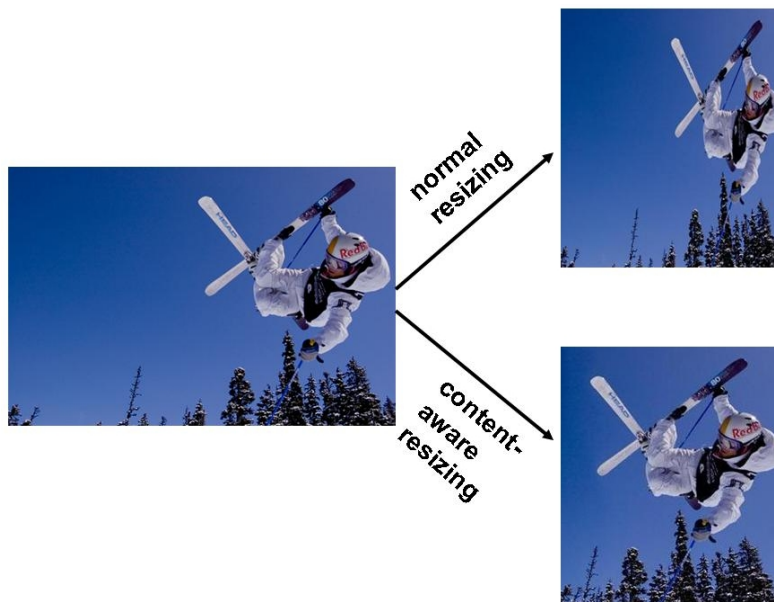
I. Short answer problems [10 points each]

1. Consider the image to the right. Say we want to filter this image so that we get the maximal responses at the centers of the “holes”. Describe in detail how to generate an appropriate filter.



2. Write pseudo-code (English, not Matlab) outlining how to generate a histogram that shows the distribution of gradient orientations within an image. Assume each pixel will increment some bin of the histogram.

II. Programming problem: content-aware image resizing [80 points]



For this exercise, you will implement a version of the content-aware image resizing technique described in Shai Avidan and Ariel Shamir’s SIGGRAPH 2007 paper, “Seam Carving for Content-Aware Image Resizing”. The paper is available here: <http://www.seamcarving.com/arik/imret.pdf>. The goal is to implement the method, and then examine and explain its performance on different kinds of input images.

First read through the paper, with emphasis on sections 3, 4.1, and 4.3. **Note:** choosing the next pixel to add one at a time in a greedy manner will give sub-optimal seams; the dynamic programming solution ensures the best seam (constrained by 8-connectedness) is computed. Use the dynamic programming solution as given in the paper.

Write Matlab code with functions that can do the following tasks:

- Compute the energy function at each pixel using the magnitude of the x and y gradients (equation 1 in the paper)
- Compute the optimal vertical seam given an image
- Compute the optimal horizontal seam given an image
- Reduce the image size by a specified amount in one dimension (width or height decrease)
- Display the selected seam on top of an image
- Functions with the following interface:

```
[output] = reduceWidth(im, numPixels)
```

```
[output] = reduceHeight(im, numPixels)
```

These functions take an input image `im`, and a parameter specifying how many seams to carve, from the width or height, respectively. The image `im` will be a $h \times w \times 3$ `uint8` matrix, which is what `imread` returns for a color image. Put these functions in file named `reduceWidth.m` and `reduceHeight.m`

Set up scripts so that you can play with the seam removal and specify different combinations of horizontal and vertical removals. Apply your system to the provided images. View the results in color, but note that the gradients should be computed with the grayscale converted image.

Matlab hints:

- Useful functions: `imfilter`, `fspecial`, `imread`, `imresize`, `rgb2gray`, `imagesc`, `imshow`, `subplot`;
- To plot points on top of a displayed image, use `"imshow(im);"` followed by `"hold on;"` followed by `"plot(...)"`.
- Be careful with `double` and `uint8` conversions as you go between computations with the images and displaying them – filtering should be done with doubles.

Answer each of the following, and include image displays where appropriate:

1. [10 points] Run your `reduceWidth` function on the provided `seals.jpg` with `numPixels = 325` (in other words, shrink the width by 325 pixels). Display the output.
2. [10 points] Display (a) the energy function output (total gradient magnitudes $e_1(I)$) for the provided image `seals.jpg`, and (b) the two corresponding cumulative minimum energy maps (M) for the seams in each direction (use the `imagesc` function). Explain why these outputs look the way they do given the original image's content.
3. [10 points] For the same image `seals.jpg`, display the original image together with (a) the first selected horizontal seam and (b) the first selected vertical seam. Explain why these are the optimal seams for this image.

4. [10 points] Make some change to the gradient computation (i.e., filter used and/or its parameters) and explain the impact on the results for some example.
5. [40 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.

Include results for the **three provided images**, plus at least **three images of your own** choosing. Include an example or two of a “bad” outcome. Be creative in the images you choose, and in the amount of combined vertical and horizontal carvings you apply. Try to predict types of images where you might see something interesting happen. It’s ok to fiddle with the parameters (seam sequence, number of seams, etc) to look for interesting and explainable outcomes.

For each result, include (a) the original input image, (b) your system’s resized image, (c) the result one would get if instead a simple resampling were used (via Matlab’s `imresize`), (d) the input and output image dimensions, (e) the sequence of enlargements and removals that were used, and (f) a qualitative explanation of what we’re seeing in the output.

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

Below are ways to expand on the system you built above. If you choose to do any of these (or design your own extension) include in your writeup an explanation of the extension as well as images displaying the results and a short explanation of the outcomes. Also include a line or two of instructions telling what needs to be done to execute that part of your code.

1. Allow a user to mark an object to be removed, and then remove seams until all pixels on that object are gone (as suggested in section 4.6 of the paper). Either hard-code the region specific to the image, or allow interactive choices (Matlab’s `ginput` or `impoly` functions are useful to get mouse clicks or draw polygons).
2. Design an alternate energy function, instead of the gradient magnitude. Explain your choice, and show how it can influence the results as compared to using the gradient magnitude. Choose an image or two that illustrates the differences well.
3. To avoid warping regions containing people’s faces, have the system try to detect skin-colored pixels, and let that affect the energy map. Try using the hue (H) channel of HSV color space (see Matlab’s `rgb2hsv` function to map to HSV color space). Think about how to translate those values into energy function scores.
4. Optimize for the order of horizontal and vertical seams, as described in Section 4.2 of the Avidan & Shamir paper. Show an example where this compares favorably to hard-coding the sequence.

Submission instructions: what to hand in

Electronically:

- Your documented Matlab code (including functions `reduceWidth.m` and `reduceHeight.m`)
- A first pdf file named `file1.pdf` containing the following:
 - Your name at the top
 - Your answers to Part I, numbered.
 - A brief explanation of your implementation strategy: a short paragraph or two describing in English what you have computed.
 - Your responses and image results for questions 1 through 4 in Part II, numbered. Insert image figures in the appropriate places for these questions.
- A second pdf file named `file2.pdf` containing the following:
 - Your name at the top
 - Your responses and image results for question 5 in Part II.
 - (optional): any results and descriptions for extra credit portions.
 - This file will be posted online – be sure to credit any photo sources.

Submit all the above with one call to turnin:

```
>> turnin --submit jaechul pset1 file1.pdf file2.pdf  
reduceWidth.m reduceHeight.m <otherFunction.m> ...  
<otherFunction.m> etc.
```

Hardcopy:

- Printout of the two pdf files above, and bring to class on Tuesday 9/22/09. The hardcopy must be identical to what is submitted electronically by Monday night.

Tip: How to save as pdf? If you have a pdf printer installed on your computer, you can convert a document prepared in Word to pdf. The CS machines have `openoffice` installed, which will also allow you to Save as... a pdf file. Or, if you work in Latex, you can use `pdflatex`, or compile to a ps and then convert.

Image acknowledgements: Thanks to the following Flickr users for sharing their photos under the Creative Commons license:

`seals.jpg` is provided by allotrope.
`trees.jpg` is provided by russelljsmith.
`groceries.jpg` is provided by The Consumerist.