Homework 2: Image Filtering and Hybrid Images

Brief

Due date: Feb 17, Tuesday, end of day
Team work: encouraged, maximum number of members, 3
Totals points: 100+20pts

Overview

The goal of this project is to write an image filtering function and use it to create hybrid images using as a simplified version of the SIGGRAPH 2006 paper by Oliva, Torralba, and Schyns. *Hybrid images* are static images that change in interpretation as a function of the viewing distance. The basic idea is that high frequency tends to dominate perception when it is available, but, at a distance, only the low frequency (smooth) part of the signal can be seen. By blending the high frequency portion of one image with the low-frequency portion of another, you get a hybrid image that leads to different interpretations at different distances.

Details

This project is intended to familiarize you with Numpy and image filtering. Once you have created an image filtering function, it is relatively straightforward to construct hybrid images. If you don't know Numpy, you might find the online tutorials on Numpy useful.

Image Filtering. Image filtering (or convolution) is a fundamental image processing tool. See chapter 3.2 of Szeliski and the lecture materials to learn about image filtering (specifically linear filtering). Python/Numpy/Scipy have numerous built in and efficient functions to perform image filtering, but you will be writing YOUR OWN such function from scratch for this assignment. More specifically, you will implement `my_imfilter()` which imitates the default behavior of the build in scipy.`ndimage.filter` module and the MATLAB imfilter.m function. As specified in `my_imfilter.py`, your filtering algorithm must (1) support grayscale and color images (2) support arbitrary shaped filters, as long as both dimensions are odd (e.g. 7x9 filters but not 4x5 filters) (3) pad the input image with zeros or reflected image content and (4) return a filtered image which is the same resolution as the input image.

Hybrid Images. A hybrid image is the sum of a low-pass filtered version of the one image and a high-pass filtered version of a second image (Hint: how do you get a high pass image? Remember unsharp mask?). There is a free parameter, which can be tuned for each image pair, which controls *how much* high frequency to remove from the first image and how much low frequency to leave in the second image. This is called the "cutoff-frequency". In the paper it is suggested to use two cutoff

frequencies (one tuned for each image) and you are free to try that, as well. In the starter code, the cutoff frequency is controlled by changing the standard deviation of the Gausian filter used in constructing the hybrid images. Plot the frequency cut off for both images.

We provide you with five pairs of aligned image, which can be merged reasonably well into hybrid images. The alignment is important because it affects the perceptual grouping (read the paper for details). We encourage you to create additional examples (e.g. change of expression, morph between different objects, change over time, etc.). See the **hybrid images project page** for some inspiration.

Forbidden functions you can use for testing, but not in your final code: `ndimage.filters, ndimage.convolve, ndimage.convolve1d`

Bells & Whistles (Extra Points)

Extra points will only be added if you successfully finished the above steps.

You can do the image filtering in Fourier domain.

You can use novel images other than the 5 pairs.

Provide Fourier transform images of the high pass and low pass images.

In addition, can you allow the user to adjust the weights? Can you create user-interface on your web?

Web-Publishing Results

All the results for each project will be put on the course website so that the students can see each other's results. The professor will select "winning" projects that impress us and there will be in class presentations for these projects. If you do not want your results published to the web, you can choose to opt out. If you want to opt out, email me.

Write up

For this project, and all other projects for the rest of the semester, you must do a project report in HTML. We provide you with a placeholder .html document, which you can edit. In the report you will describe your algorithm and any decisions you made to write your algorithm a particular way. Then you will show and discuss the results of your algorithm. In the case of this project, show the results of your filtering algorithm and show some of the intermediate images in the hybrid image pipeline (e.g. the low and high frequency images). Also, discuss anything extra you did. Feel free to add any other information you feel is relevant. If you performed this

assignment as a team, please indicate this in the start of your webpage and credit each person for his/her contributions.

The details of how to setup your webpage on the AU server can be found here:

http://www.american.edu/oit/network/WebPage-Students.cfm

But you don't have to have Internet to view the webpage.

Handing in

This is very important, as you will lose points if you do not follow instructions. Every time after the first that you do not follow instructions, you will lose 5 points. The folder you hand in must contain the following:

- README - text file containing anything about the project that you want to tell the TAs
- code/ - directory containing all your code for this assignment
- html/ - directory containing all your html report for this assignment (including images). Only this folder will be published to the course web page, so your webpage cannot contain pointers to images in other folders of your hand-in.
- html/index.html - home page for your results
- Zip the folder and submit to blackboard.
- You can upload your folder onto your own website. If you do so, you are still required to submit the folder and the address of the website via blackboard.

Rubric

- +50 pts: Working implementation of image filtering in `my_imfilter.py`
- +30 pts: Working hybrid image generation
- +20 pts: Writeup with several examples of hybrid images
- +20 pts: Extra credit (up to ten points)
- -5*n pts: Lose 5 points for every time (after the first) you do not follow the instructions for the hand in format
- Late policy applies!

This assignment was modified from multiple resources (James Hayes at Brown, and Derek Hoiem at UIUC, etc). But web plagiarism is strictly forbidden.