CS4670 Spring 2015

HW1 - Filtering, Sampling, Features

Due: ~~Feb 24 (Tue)~~ Feb 26 (Thu), 2015
Homeworks to be done alone.
Remember the academic integrity policies.
Consult piazza for clarifications on the HW1 FAQ.

1 Filtering

1. You have a floating point image of dimension $m \times n$ and you want to convolve it with a kernel of size $p \times q$, where $m \geq p$ and $n \geq q$. Assume that we use zero-padding for convolution on the edges. How many floating point addition and multiplication operations are needed?

2. Separable kernels. You learned in the lectures that some 2D kernels can be linearly separated into two 1D kernels, which lets you perform convolution in a much faster way. Let us assume the 2D kernel above can be decomposed into 1D kernels of size $1 \times q$ and $p \times 1$. How many addition and multiplication operations will then be necessary to perform the convolution?

2 Features and Robust Model Fitting

Consider two windows in an image.

Window A

1	1	1	1
1	1	1	1
2	2	2	2
2	2	2	2

Window B

1	1	1	1
1	1	1	1
1	1	2	2
1	1	2	2

1. Given the above windows of an image, compute the structure tensor used by the Harris corner detector. In this case, the structure tensor refers to the sum of the 2×2 second moment matrices of all the pixels in the window. When computing the gradient, assume that pixels outside the window are equal to the adjacent pixel inside the window.

2. Compute the eigenvalues for each of these two matrices. Recall that the eigenvalues of a 2x2 matrix M can be computed using the trace $T = M_{11} + M_{22}$ and determinant $D = M_{11}M_{22} - M_{12}M_{21}$.

3. Compute the Harris score for each window.

4. Which window has a corner based on its Harris score? Describe (briefly) how you make this determination.

3 RANSAC

1. Suppose we are robustly fitting a circle to a set of 2D points using RANSAC. In each iteration, how many points do we choose for our initial hypothesis set?
2. Now suppose we want to robustly fit a circle (not a sphere) to a set of 3D points. How many points are chosen for a hypothesis set?
3. Suppose we have a model (for example, general conic sections) for which a minimal hypothesis set is 5 points. If we have 10,000 points and we expect that 85% of these are inliers, how many iterations of RANSAC do we need to run in order to find the correct fit with probability 0.95?
4. Now suppose 85% of the points are outliers. How many iterations are required now?

4 Downsampling and Aliasing

1. What is the maximum frequency sine wave that can be represented in a 1D discrete image F with N pixels? Write your answer in the form $F(x) = \sin(\omega_{max}\, x + \pi/2)$ where $x = 0, 1, \ldots, N-1$. (Hint: you may find it useful to plot F).
2. Although we have established that ω_{max} is the maximum representable frequency, we could still try and construct an image from a sinusoid with a frequency higher than max. If we construct a 1D discrete image from the formula $F(x) = \sin(1.5\, \omega_{max}\, x + \pi/2)$, what frequency will we observe when looking at F?
3. Before subsampling an image, prefiltering is used to avoid aliasing. What property must the prefilter have in order to avoid aliasing, when subsampling an arbitrary 1D image by ratio $s > 1$?

5 Upsampling

To upsample a low resolution 1D image F with upsampling rate r > 1, we resample the input F using an interpolation kernel h to get a higher resolution output G, where $G(i) = \sum_k F(k)h(i - rk)$. The interpolation kernel h is typically chosen to have the property $\sum_k h(i - rk) = 1$ for all i.

1. What happens to the output if we chose a kernel with the property $\sum_k h(i-rk) < 1$?
2. What about if $\sum_k h(i - rk) > 1$?
3. If we handle the boundary problem by extending F with zeros, what will happen to the boundary of G?

6 Unconventional interpolation

Consider a 1D image F with 3 pixels: $F(0) = 2$, $F(1) = 1$, $F(2) = 1$.

1. Convert the discrete image F to a continuous function f and plot it (Hint: each pixel becomes an impulse).

2. We are going to use an unconventional upsampling kernel h that is half triangle filter, half box filter:

$$h(x) = \begin{cases} x+1 & \text{if } x \in (-1, 0] \\ 1 & \text{if } x \in (0, 1/2] \\ 0 & \text{else} \end{cases}$$

   Plot h. Indicate discontinuities using closed and open circles.

3. Plot the result of convolving f and h; call this function g. Indicate discontinuities using closed and open circles.

4. Resample g at twice the resolution to get a discrete image G (sample from 0 to 2.5 to get a 6-pixel image).

7 Multi-scale pyramids

Consider the Laplacian Pyramid from Lecture 04, slide 37.

1. Assume that we have only one level of the pyramid ($n = 1$). Write down an expression that reconstructs the input (G_0) given the Laplacian pyramid representation L_0, G_1.

2. Assume that we have 3 levels ($n = 3$). Do the same for the Laplacian pyramid representation L_0, L_1, L_2, G_3.

3. Is any information lost as part of this reconstruction?

4. Is it possible to always reconstruct the input (G_0) given only the Gaussian Pyramid levels G_1, G_2, G_3? If so, write down the expression that performs this reconstruction.