GIF-4105/7105 Computational Photography
Winter 2015
Midsemester exam
February 26, 2014
Total time: 110 minutes

This exam has 7 questions on 8 pages (including this one), and is worth 20% of the total grade for the semester. Make sure you have all the pages. The following rules apply:

- You are allowed a handwritten double-sided 8.5×11 helper sheet;
- **Write detailed answers**. That way, you can have points for partial answers;
- Write your answers in the accompanying blue booklet, and hand in the booklet and this questionnaire.

The following table indicates the points distribution for each question.

Question:	1	2	3	4	5	6	7	Total
Points:	10	10	20	30	10	10	10	100

The exam also has additional questions.

- *Undergraduates (4105)*: you can answer the questions to get bonus points;
- *Graduates (7105)*: the questions are mandatory, and your total score will be computed over 130 points (100 base points + 30 additional points).

The following table indicates the points distribution for each additional question.

Question:	1	2	3	4	5	6	7	Total
Points:	5	0	0	10	5	0	10	30

Good luck!

1. **Pixels and colors (10 points, +5 for graduate students)**

   (a) (2 points) What is an image histogram?

   (b) (3 points) Name the algorithm that uses the histogram of an image to improve its contrast, and describe it.

   (c) (3 points) Name one advantage and disadvantage for each one of the following color spaces:

      i. RGB;

      ii. HSV;

      iii. LAB.

   (d) (2 points) In compression, we often split an image into its luminance (intensity) and chrominance (color). Which one would you compress more aggressively? Why?

   (e) (5 points, **Additional question**) In the figure below, how would you adapt the algorithm from question (b) to match the histogram of image 1 to that of image 2? The algorithm should result in image 1*, which has a histogram that is similar to image 2.

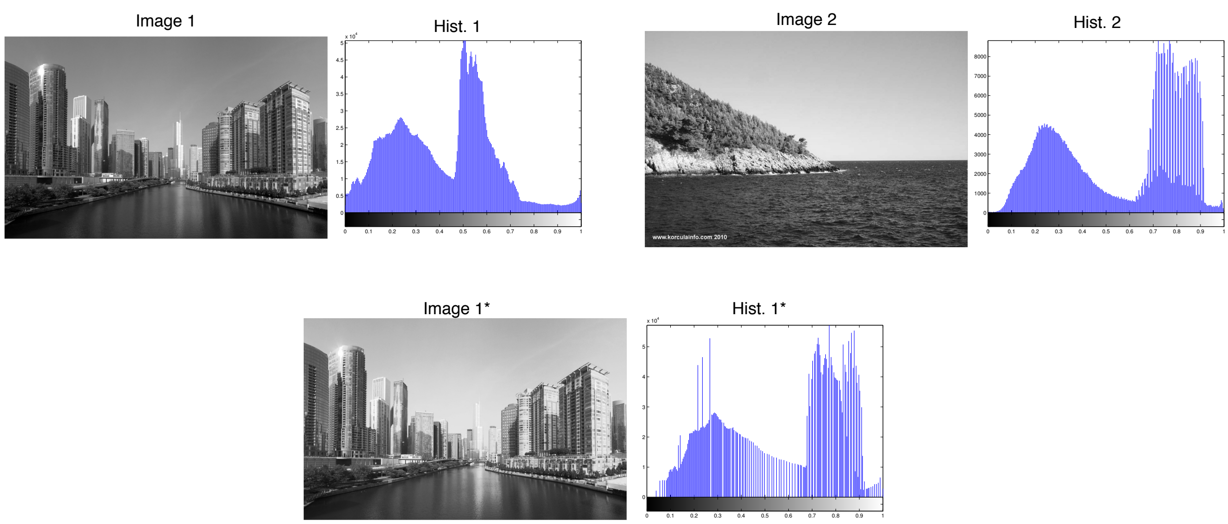

2. **Spatial domain filtering (10 points, +0 for graduate students)**

The 1-D signal in the figure below (in bold) is corrupted by noise such that its sampled values are the following: 1, 3, 1, 2, 3, 5, 7, 8, 9, 8. In this question, we would like to detect the edge of the original signal.

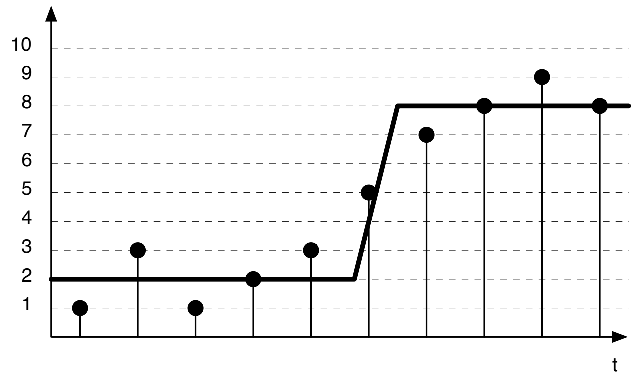

(a) (2 points) The gradient $g(t)$ of a signal $f(t)$ can be computed by $g(t) = f(t+1) - f(t)$. Write the 1-D filter that corresponds to this operation.

(b) (3 points) Apply this filter to the sampled 1-D signal above to compute $g(t)$.

(c) (5 points) How can we detect the edge from the sampled signal?

3. **Spectral domain filtering (20 points, +0 for graduate students)**

   (a) (4 points) Associate the FFT to the corresponding image.


1


2


3


4

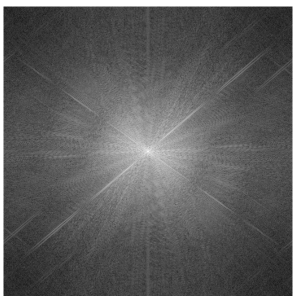
A

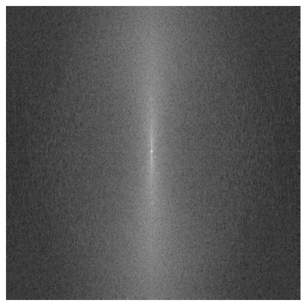
B

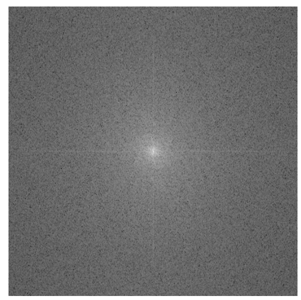
C

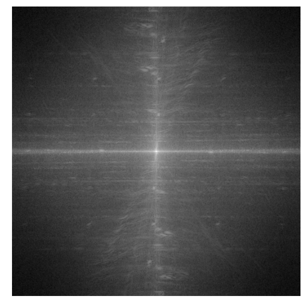
D

   (b) (5 points) Sometimes, clothing creates weird effects when seen on TV (see image below, for example). What is this phenomenon, and why does it happen?

   (c) (5 points) In theory, what are the two strategies to reduce this undesired effect?

   (d) (6 points) Filtering is often used to smooth images. Would you prefer a box filter, or a gaussian filter? Explain why by using the convolution theorem.

4. **Transformations and morphing ()**

(a) (5 points) A point $p = (x, y)$ undergoes the following transformations (in order): a translation of $(2, -7)$, followed by a scale factor of $(2, 3)$, followed by a rotation of $-90°$. Knowing that the 2×2 matrix representing a rotation is

$$R = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix},$$

what is the resulting matrix? Detail your work.

(b) The quadrilateral $ABCD$ is deformed by an affine transformation T, and generates the quadrilateral $A'B'C'D'$, as shown below.

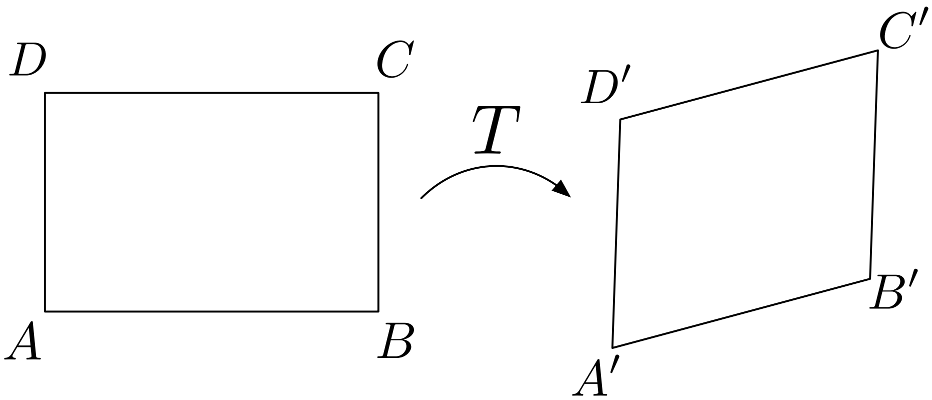

i. (5 points) What is the minimum number of correspondences needed to compute the parameters of T? Explain why.

ii. (10 points) How can we compute these parameters from the correspondences?

(c) (5 points) We give you an image and a transformation matrix M. Describe the procedure to warp the image according to M.

(d) (5 points) The cat and the puma are both felines, so why isn't a global transformation not sufficient to warp the left image into the right one? What should we use instead?

(e) (10 points, **Additional question**) To morph an image into another, the approach we suggest you employ in HW3 is to use a triangulation, and to warp each triangle towards its corresponding one. What is the problem with this approach? Describe an alternative approach that doesn't have this problem.

5. **Removing an object (10 points, +5 for graduate students)**

   Removing an object is a common image editing operation, and many algorithms can be employed for that purpose. For each photo below, describe solutions that you would use to remove the indicated object. *You must use a different strategy for each photo.*

   (a) (5 points) Graffiti on the wall:

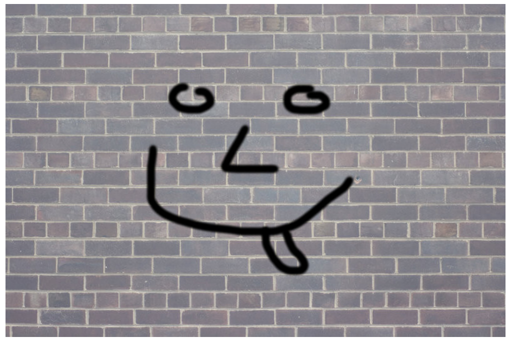

   (b) (5 points) The lady in the middle of the image:

   (c) (5 points, **Additional question**) The bird in the sky (note: to get the points, you must propose a solution different than the ones you already used above)::

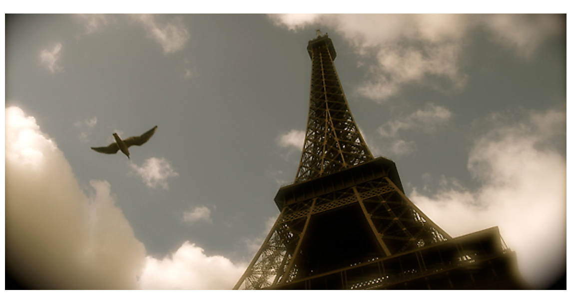

6. **Cutting (10 points, +0 for graduate students)**

   (a) (2 points) How can we compute the lowest cost path from one side of an image to the other, without having to enumerate all possible paths?

   (b) (5 points) Compute the minimum cost path that goes from the left to the right of the following image by using this strategy.

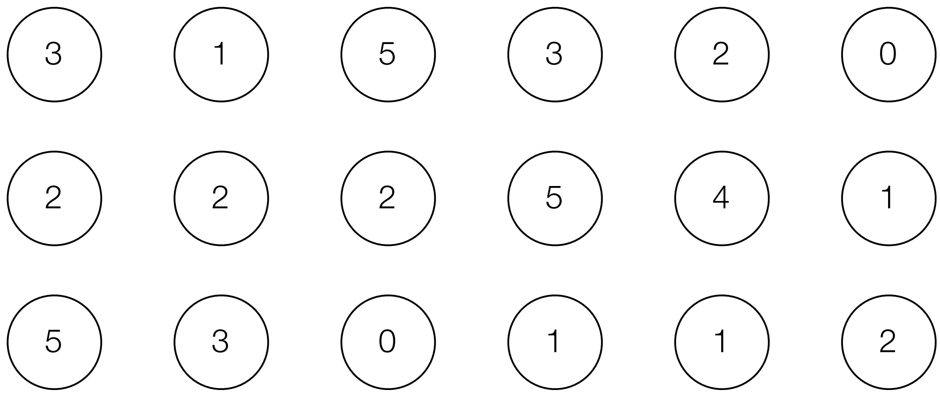

   (c) (3 points) In the “intelligent scissors” algorithm, what would happen if we were to replace the boundary cost with a value of 0 for each pixel?

7. **Image compositing (10 points, +10 for graduate students)**

   (a) (5 points) I would like to insert the bear in the image below, but I'm not sure about which algorithm to use. Help me by describing the problems that could arise by using the following algorithms:

      i. copy-paste;

      ii. feathering;

      iii. gradient domain blending.

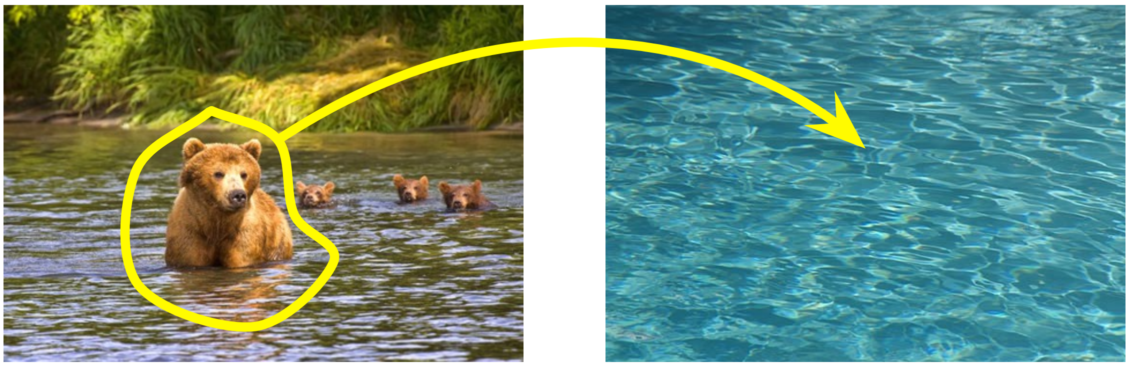

   (b) (5 points) After much reflection, I decide to use the gradient domain blending algorithm. How should I combine the gradients coming from the two images?

   (c) (10 points, **Additional question**) Gradients can also be used for operations other than image compositing. For example, by computing the gradients, modifying them, and re-integrating, I can generate the effect shown in the following image. Note how the main regions of the image are preserved, but the small-scale details (like textures) are lost. Describe an approach to synthesize such a result by editing gradients.

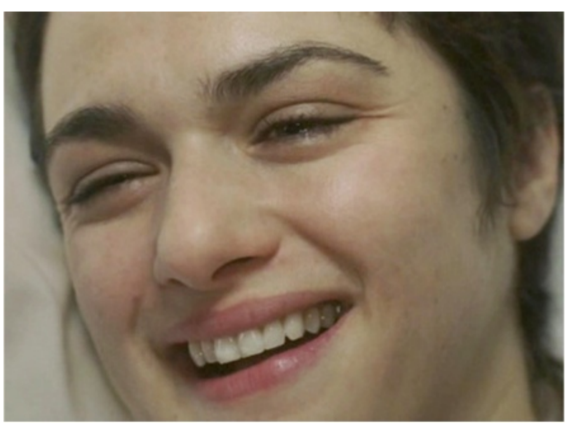

Original image

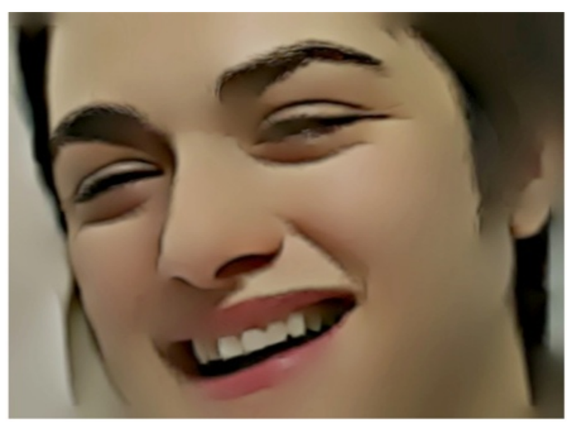

Non-photorealistic rendering