Computer Vision

Professors : James L. Crowley, Edmond Boyer

M2R GVR | Mid-term Exam, November 2010

Test conditions: All documents and reference materials are authorized. You may NOT communicate with anyone other than the exam Proctor or the course professor. You must answer all questions in INK on the official exam paper. You may use scratch paper to prepare your answer, but your scratch paper will not be graded. You may respond in English or French (or both), but you MUST write legibly. Use mathematics as well as English and/or French to communicate. Illegible text will not be graded.

Duration : Maximum 3 hours.

For an image *P(i,j)*, the gradient $\vec{\nabla}P(i,j)$ is a vector composed of the first derivatives in the row and column directions. In the following questions, assume that *P(i,j)* is a luminance (black and white) image of size 1024 x 1024 pixels with 8 bits per pixel.

1) You are asked to compute the gradient $\vec{\nabla}P(i,j)$ using convolution with sampled Gaussian derivatives, $\vec{\nabla}G(i,j)$. Give the formulae for the sampled Gaussian derivatives as well as the formulae for the 2D convolution.

2) (2 points) What is the minimum size support window that can be used for a Gaussian derivative with σ=2? What is the computational cost for such a convolution in terms of additions and multiplications when implemented as a 2-D convolution?

3) (2 points) Show that the sampled Gaussian derivatives are separable. That is, that convolution with the sampled Gaussian derivative can be implemented as a sequence of convolutions with 1-D filters in the row and column directions. What is the computational cost in terms of additions and multiplications for the convolution of the image with sampled Gaussian derivatives at σ=2 when implemented as convolution with separable 1D components?

4) (2 points) Show that a 1D Gaussian low pass filter with σ=2 can be implemented as a series of convolutions with 1D Gaussian low pass filter with σ=1. How many convolutions are needed to compute a 1D Gaussian low pass filter with σ=2 as a series of convolutions with a 1D Gaussian low pass filter with σ=1? What is the computational cost of this series of convolutions?

5) Is it possible to implement the convolution with a 1-D Gaussian Derivative filter with σ=2 as a series of convolutions with 1-D Gaussian filters? If yes, how? If no, why not?

6) Given the Gradient of the image $\vec{\nabla}P(i,j)$ in the row and column directions, give a formulae to determine the gradient at pixel (i,j) in an arbitrary direction θ.

7) Given the Gradient of the image $\vec{\nabla}P(i,j)$, give the formulae to determine the direction of maximum gradient, θ_{max}(i,j) at each pixel i,j.

8) Give a formula for the coefficients for a binomial low pass filter at σ=2.

9) Show that a binomial low pass filter at σ=2 can be implemented as a series of convolutions with the binomial filter [1 2 1]. How many such convolutions are required?

10) Assume that you use [1 0 -1] as a first derivative operator. What is the computational cost (additions and multiplications) for computing the Gradient of an image at $\sigma=2$ using a series of convolutions with 1-D binomial filters [1 2 1] in the row and column directions?

11) (2 points) What is the computational cost to compute the gradient of the image at $\sigma=2$ using a resampled binomial pyramid algorithm with $\sqrt{2}$ resampling and [1 0 -1] as a first derivative operator?

12) (5 points)You are given an RGB image of a room illuminated with light that is predominantly green. You observe a specular reflection on the face of person in the scene. Explain how to use this specularity to correct the color balance of the image to resemble an image illuminated by white light.