Computer Vision

Professor: James L. Crowley

M2R GVR Mid-term Exam November 2014
Duration: 3 hours

Test conditions: All documents and reference materials are authorized. You may NOT communicate with anyone other than the exam Proctor or the course professor (or his representative). 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. Illegible text will not be graded. Use mathematics as well as English and/or French to communicate.

1) (2 points) What is a Lateral Geniculate Nucleus (LGN)? What function does it provide for the human visual system?

2) (2 points) What is the Superior Colliculus? What function does it provide for the human visual system?

3) (16 points) A ratio of histograms can be used with a vector of image derivatives, $\vec{X} = (P_x, P_y, P_{xx}, P_{yy}, P_{xy})$ to build a pixel level Bayesian detector for patterns. In this problem, we will use this approach to build a detector for human faces in gray-scale images. Assume that you are given a training set of M images, $P_m(x,y)$, with face regions ($P(x,y) \in face$) indicated in each image, where each image is of size 1024 x 1024 pixels, and each face is of (approximately) of size 64x64 pixels.

a) (2 points) Give the full formulae for the sampled Gaussian derivatives for each of the filters $(G_x, G_y, G_{xx}, G_{yy}, G_{xy})$.
b) Give the formula for the 2-D convolution $P_x(x,y) = G_x * P(x,y)$.
c) Explain how to use the training data to obtain the histogram $h(\vec{X})$ for estimating the probability of a vector $p(\vec{X})$.
d) Explain how to use the training data to train the histogram $h_{face}(\vec{X})$ for calculating the probability of a vector given that the pixel is part of the face, $p(\vec{X} \mid P(x,y) \in face)$.
e) (2 points) How large should the histograms be?
f) How can you estimate the a-priori probability that a pixel will be part of a face $p(P(x,y) \in face)$?
g) (2 points) Explain how to use the histograms to determine the probability that a pixel represents part of a face given the value of its derivative vector, $p(P(x,y) \in face \mid \vec{X})$.
h) (2 points) How can you determine what value of σ to use?
i) Explain how to compute an estimate of the likelihood that an image window of size WxH contains a face based on the estimated probabilities at each pixel.
j) Explain how to compute the center of gravity of the face region from the probability of face in the image window of size WxH. The center of gravity should be expressed in image coordinates.
k) Explain how to compute the second moment of the face region.
l) Explain how to determine the width, length and orientation of the face region from the second moment.