Computer Vision

Professor: James L. Crowley

M2R GVR Mid-term Exam 1 December 2016

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) (4 points). An RGB camera observes a light gray table from a distance of 2 meters and an angle of 45 degrees (relative to the normal of the table. The center of the image is the center of the table. The table is illuminated by an LED light source that is located directly overhead the table. There are no other sources of light in the room. When observed by the camera, the light source appears as a disk with relative RGB colors of (1, 1, 0.5) A sphere is placed in the middle of the table, such that it appears in the middle of the image. The albedo of the sphere is a mixture of lambertian and specular reflections with a coefficient of 0.10% specular and 0.90% lambertian. The sphere is coated with a pigment that absorbs 50% of the red photons, 100% of the green photons and none of blue photons.

a) What are the relative RGB values observed for the specular reflection on the sphere?

b) What color does the specular reflection appear to have for a human observer?

c) What is the RGB color of the sphere outside of the specular reflection?

d) What color does the sphere appear to have for a human observer?

2) (6 points) Assume a cube with 6 visible corners. Explain how to compute the projection matrix M_s^i for the camera from the image position of the corners of a cube. What is the minimal number of corners required? How would you use ALL the corners? Which gives a more precise answer? Why?

3) (8 points) Assume a 1024 x 1024 gray-scale image *P(i,j)* at 8 bits per pixel and its Gaussian Pyramid *P(i,j,k)*.

a) Give the full formula for a 2-D Gaussian derivative filter $G_x(i,j;\sigma_k)$.

b) Give the formula for the 2-D convolution $P(i,j,k) = P * G_x(i,j;\sigma_k)$

c) What is the minimum spatial extent (size in pixels) for the filter $G(i,j;\sigma_k)$ for $\sigma_k = 1, 2,$ and 4?

d) How would you assure that the sum of the coefficients is 1 for the filter $G(i,j;\sigma_k)$?

e) Show how to compute *P(i,j,k)* for $\sigma_k = 4$ as a series of 1D convolutions with the 1D Gaussian filters with $\sigma_0 = 1$.

f) Suppose that $\sigma_k = 4$. What sample distance, Δx_k, we can use to resample $G(i,j;\sigma_k)$ with minimal loss to aliasing?

g) Show that $P_x(i,j,k) \approx P(i+1,j,k) - P(i-1,j,k)$

h) Show that $\nabla^2 P(i,j,k) = P_{xx}(i,j,k) + P_{yy}(i,j,k) \approx P(i\Delta x_k, j\Delta x_k, k) - P(i\Delta x_{k-1}, j\Delta x_{k-1}, k-1)$

4) (2 points) A colleague suggests that you can calibrate a projective camera model from the image position of the corners of a checkerboard placed on the table. Is he correct? If yes, give the equations and algorithm for the calibration. If he is not correct, explain why.