2004 Paper 9 Question 11

Computer Vision

(a) Why is the performance of current face recognition algorithms so poor? Address both what makes the problem domain intrinsically so challenging, and the shortcomings of the strategies adopted in the design of current algorithms. Comment on possible directions for improving performance. [8 marks]

(b) For what size of filter kernel does it become more efficient to perform convolutions by instead computing Fourier Transforms, and why? [2 marks]

(c) For an aligned stereo pair of cameras separated by base distance b, each with focal length f, when a target point projects outside the central axis of the two cameras by amounts α and β:

(i) What is the computed target depth d? [2 marks]

(ii) Why is camera calibration so important for stereo vision computations? [1 mark]

(iii) Identify *four* relevant camera degrees-of-freedom and briefly explain their importance for stereo vision algorithms. [2 marks]

(d) What does the Spectral Co-Planarity Theorem assert about translational visual motion, and how the parameters of such motion can be extracted? [2 marks]

(e) What information about the shape and orientation of an object can be inferred, and how, from the extraction of texture descriptors; and what is the role of prior assumptions in making such inferences? [3 marks]