


5 Computer Vision (jgd1000)

- (a) (i) For an image $I(x, y)$, define its gradient vector field $\vec{\nabla}I(x, y)$. [1 mark]
- (ii) Why is this vector field a useful thing to compute? [1 mark]
- (iii) Define the gradient magnitude over the image plane (x, y) . [1 mark]
- (iv) Define the gradient direction over the image plane (x, y) . [1 mark]
- (v) Explain how the gradient vector field is used in the Canny edge detector, the main steps in its use, and its advantages. [3 marks]

(b) A Bayesian classifier uses observations x to assign visual objects to either one of two classes, C_1 or C_2 . Their baseline prior probabilities are $p(C_1)$ and $p(C_2)$, with sum $p(C_1) + p(C_2) = 1$. Observations x have unconditional probability $p(x)$, and the class-conditional probabilities of a given observation x are $p(x|C_1)$ and $p(x|C_2)$.



- (i) Using the above quantities provide an expression for $p(C_k|x)$, the likelihood of class C_k given an observation x . (Here $k \in \{1, 2\}$.) [2 marks]
- (ii) Provide a decision rule using $p(C_j|x)$ and $p(C_k|x)$ for assigning classes $\{C_j, C_k\}$ based on observations x , that will minimise errors. [2 marks]
- (iii) Now express your decision rule instead using only the quantities $p(C_j)$, $p(C_k)$, $p(x|C_j)$, $p(x|C_k)$, and relate it to the diagram above. [1 mark]
- (iv) If the classifier decision rule assigns class C_1 if $x \in R_1$, and C_2 if $x \in R_2$ as shown in the figure, what is the total probability of error? [2 marks]

(c) Propose an algorithm for shape classification that could correctly classify all of the objects shown here as cashew nuts, despite their variations in size (or hence distance), pose angles, colours, and intrinsic shapes. How can shape grammars, active contours, boundary descriptors, zeroes of curvature, and codon constraints enable a classifier to achieve those invariances?



[6 marks]