UG Exam Performance Feedback
Third Year
2017/2018 Semester 2

COMP37212 Computer Vision

Aphrodite Galata
Carole Twining (FLS)

Comments

Question 1:
Part a) There were some good answers to this part of the question. It was standard bookwork, but the main reasons for people not gaining marks was either not sufficient detail (like exactly WHAT is being optimised!), missing out some of the steps, or not showing that they had fully grasped some of the issues of aligning a group of shapes.
Part b) The first part, how to construct an SSM, was reasonably well done. As before, lack of detail, not including the key concepts/names, or missing some steps were the main reason for loss of marks. But there were some really excellent answers. The second part, of constructing an ASM, was not quite so well done. Although most people had remembered the difference between an ASM and an SSM, some people told me more about how it was used in search, rather than how it was built. But overall, some good answers, and it was very good to see that many people remembered and included the relevant diagrams or equations in their answers, which is a much more efficient and precise way of communicating the key concepts than lines and lines of prose, unless you write very carefully and at length.
Part c) This last part was more challenging, and although many people remembered the key idea of AAMs and texture, and the relevant amounts of image information used by ASM versus AAM, not everyone made it all the way to the combined modes of variation in an AAM. But there were many very good attempts, and a few excellent ones!

Question 2
Part a) There were some good answers to this question. Although almost no one made the mistake of taking a matrix multiplication as a convolution, many people neglected to provide enough details to make it totally clear that they understood exactly what convolution was. The use of the words equation/diagram in the question were a clue that some detail was expected.
Most people could come up with an example of a suitable filter, although explaining exactly how it did what you wanted wasn't always quite so well done. Some people told me about rank filtering instead, which wasn't what was asked for. But overall, most people knew quite a bit about what was required here, and marks were only lost for not providing enough precise detail.

Part b) Almost everyone could remember some detail about the Hough Transform, and it showed that people could remember the equations, diagrams and/or sketches from lectures! What many people missed was how you went from the edge-strength image to a set of points suitable for applying the Hough Transform. Not everyone spotted the simplification offered by ONLY looking for horizontal lines, but here it was time that was lost rather than marks. Not everyone included the accumulator array, or some idea of discrete bins being needed to count votes. But there were some excellent answers, as to how edge-strength could be included in the implementation.

Question 2.(c) : Surprisingly, performance in this question was very mixed: Some of the students that answered the question, did extremely well. Others did not give enough detail in their answers. Finally, quite few students failed to answer this question even though it examines standard bookwork!

Question 2.(d) The majority of the students did well in this question - the clue was to understand what an epipole is. Some students got full marks for this question

Question 3.(a) Generally it was well answered, however some students failed to explain how to combine an LoG with the harris corner detector to find scale invariant corner features. Here a definition (and/or drawing) of the scale-space and how you search in the scale-space was expected. Again, this question tests standard bookwork and some students got full marks.

Question 3.(b) Generally well answered.

Question 3.(c) This question was meant to check more in-depth knowledge and understanding. A number of students got full marks here however the majority failed to notice that although in the SIFT feature description that was presented in the lectures, the dimensionality of the vector space was 128, it does not always have to be so. For example, instead of having 8 bins for the histogram of gradient directions, you could choose a different number of bins for the histogram. This will change the dimensionality of the feature vector. So, in conclusion, the dimensionality parameter of the feature vector determines the _distinctiveness_ of the feature.