UG Exam Performance Feedback
Second Year
2016/2017 Semester 2

COMP27112	Computer Graphics and Image Processing	Toby Howard Tim Morris

Comments

For feedback for Section B, Q1 - Q11 (Toby Howard) see the attached report.

For feedback for Section B, Q11 - Q18 (Tim Morris)see the attached report.

COMP27112 2016-17 General exam feedback

Questions 1-11, Toby Howard.

1. Why is a double buffer necessary in a graphics rendering system? (1 mark)

This is a core concept, but about 20% of the class seemed to have no idea what this meant.

2. For 3D transformations, why do we use 4x4 matrices instead of 3x3 matrices? (1 mark)

Mostly well-answered, but many people answered too briefly, and it was not clear if they really understood.

3. What is meant by the duality of modelling and viewing? (1 mark)

Almost everyone seemed to understand this.

4. In the context of perspective projection, what does the following matrix represent? (1 mark)

About half the class knew; the others didn't, and there were some wild guesses, all wrong.

5. What is meant by projection normalization and why is it used? (2 marks)

Again, about half the class knew; the others didn't, and there were some wild guesses, all wrong.

6. Figure 1 shows part of a screen displaying a grey triangle ABC and a red triangle DEF. ABC is scan-converted first. DEF is scan-converted second, and is closer to the camera than ABC. Referring to the 2 pixels P and Q, explain how the scan conversion process ensures that ABC and DEF are displayed correctly. (3 marks)

About ¼ of the class did not discuss the z-buffer – this surprised me. Another ¼ answered the question without making any reference to P and Q, instead talking in general terms, so that's an exam technique issue.

7. What is polygon soup? Discuss whether you think it is a good or bad idea. (2 marks)

This question was generally well-answered (it should be, because it is easy).

8. Given a triangle ABC, explain how to compute the triangle's normal vector. (2 marks)

About ¼ of the class had no idea about this; this surprised me.

9. What is the fundamental difference between local and global illumination? (1 mark)

Most people knew the simple distinction, but in some case they didn't explain clearly.

10. What is the purpose of the expression ks(Rˆ · Vˆ)n in a simple local illumination model? In your answer explain the roles of each term and their ranges of values. (4 marks)

Quite well –answered. Most people knew this is an expression to estimate 'specular reflection'. Some people talked about 'spectral reflection' (never heard of that); others simply did not follow the instructions, such as not giving the range of values, as asked.

11. Describe a method for making a surface look bumpy that does not involve changing the actual geometry of the surface. (2 marks)

Most people were able to talk about 'changing the normal vector during rendering', which is the key idea. About ¼ of the class didn't seem to know about this, and talked about 'applying bumpy textures', which is not the idea here at all.

COMP27112 2016-17
General Exam Feedback
Questions 12 – 18, Tim Morris

12. Thresholding is intended to separate objects from background.
Under what circumstances will this work?
What properties of an image might cause the algorithm to fail?
How could the algorithm be modified to reduce this problem?
(3 marks)
The first part was generally well answered, the second part less so.

13. Define an edge as it relates to image processing.
Use this definition to derive an operator that will highlight edges.
(3 marks)
This was straightforward bookwork, most people realised this and gave suitable definitions.

14. Figure 2 shows an image of the Earth taken from a polar orbiting satellite. It shows land (part of North Africa at the bottom), sea and cloud. How would you process this to identify the pixels that are land, sea and cloud and hence estimate the amount and thickness of the cloud cover?
(5 marks)
People who realised that the answer should involve some form of thresholding answered this well. Other sensible answers would have received credit, but there are not any sensible alternatives.

15. Image are often corrupted by noise.
What is noise?
How would you measure the amplitude of the noise?
How would you reduce the effect of noise in an image?
(3 marks)
Similarly to question 12, the first part was well answered, the second part less so.

16. Describe a method to computing the area of a blob from a description of its outline
(2 marks)
Most people gave an account of how the chain code could be used to compute area.

17. How do you decide if two distant pixels are connected?
(2 marks)
The answer should include a definition of connectivity and similarity. Most omitted the former.

18. A central moment of area is defined by

$$M_{ab} = \sum (x - \bar{x})^a (y - \bar{y})^b f(x, y)$$

What does the 00 moment of a blob define?
How will you compute a blob's centroid?
(2 marks)
This was mostly answered correctly.