

# UG Exam Performance Feedback

## Third Year

### 2016/2017 Semester 1

---

COMP37111 Advanced Computer Graphics

Toby Howard  
Steve Pettifer

Comments Please see the report attached.

---

## Section A: Toby Howard

**Figure 1 below shows a possible triangulation of 8 points A-H. What are three undesirable properties of this triangulation, and why are they undesirable?**

This question was generally answered well, with people spotting long thin triangles, a big mixture of triangle sizes, and a gap in the convex hull, and explaining for each why they were undesirable. Some people didn't provide any explanations, so they got partial marks.

**Describe the basic principles of a technique for creating a triangulation of a point cloud that helps to eliminate the properties you described in the previous question.**

I was looking for the Delaunay method here, and about 75% of people described this. To get full marks the level of explanation should have been about the same as in the lectures; people who gave a sketchy explanation got partial marks. The remaining 25% seemed to know nothing about Delaunay, and invented their own methods; they got partial marks if they said sensible things.

**What is the 'occlusion problem' in laser scanning?**

This was an easy question, and most people were able to answer it sensibly in a few sentences, often with an example scene (good idea). Worryingly, a few people had no idea at all.

**Describe two different techniques for attempting to solve the occlusion problem.**

For full marks here I was looking for the multi-scan and zipper method; and the volume diffusion method. The question was largely well-answered, although a few people were unable to say anything sensible.

**You are building a real-time computer graphics simulation of lightning hitting a tree. Describe suitable techniques for modelling and rendering the scene.**

The key to answering this kind of question is to pause for a moment, and visualise the scene in your mind. It will look different to everyone, of course, but that's fine if the answer sensibly matches the vision, so the marking scheme was flexible - people who described their own scenario and then matched it with suitable modelling & rendering techniques got the marks. Some people seemed to have not visualised the scene at all, and listed techniques without properly relating them to the components of the scene, so I had to guess which element which technique referred to. These people therefore got partial marks.

**What is meant by the term 'calibrated camera' in the context of estimating 3D geometry from photographs?**

This was quite well-answered. For full marks I was looking for a description of intrinsic and extrinsic parameters.

**Explain how automatic feature detection is used in the estimation of 3D geometry from a video sequence.**

This question was easy to answer if you understood the concept (i.e., we calibrate between adjacent video frames using image processing algorithms (eg SIFT/SUSAN/something else sensible) to try to automatically find to a set of features that can be automatically tracked frame to frame. The changing calibrations give us an estimate of camera parameters in the sequence). If you said the above, you would get full marks. Most people did well with this question, although some talked about the different pieces of the technique, but were unable to describe how they fitted together, so they got partial marks.

## Section B: Steve Pettifer

**In classical Whitted Ray Tracing, why is it common to trace rays by starting from the eyepoint and ending at light-sources and not the other way round?**

Almost everyone got this right; those that didn't pretty much said nothing

**What is the role of shadow feelers in Whitted Ray Tracing?**

The main thing I was after here was a sense of determining direct illumination; quite a few people talked about occlusion in the sense of an object being blocked from the view, and marks were also lost for vague statements about illumination and shadows generally.

**In a scene that is lit by a single spotlight facing towards all the objects, describe what happens when a primary ray encounters a silver mirror.**

A full answer to this would require a precise description of the direction of the reflected ray; general statements about reflection typically got partial marks.

**In the same scene, what happens when a ray encounters a sphere made of green glass?**

To get all the marks you'd need to comment on the change of colour and be precise about the behaviour of the refraction effect.

**In the same scene, what happens when a ray encounters a cube covered in a matt material such as velvet?**

Several people said that because ray tracing doesn't deal with matt objects at all, that therefore the cube wouldn't get rendered; this isn't the case!

**Why is it difficult to render velvet effectively using Ray Tracing? Which alternative photo-realistic rendering technique might be used to give a better result? Briefly explain why.**

The ray-explosion that would result from attempting to render a diffuse material was generally well-described, as was radiosity as an alternative, more appropriate, solution, and most people who attempted this question got full or nearly full marks.

**Select and name one of the data structures from Figure 2 below, and explain how it could be used to accelerate the process of Ray Tracing.**

Where people selected the right data structures, the description followed nicely and got full or nearly-full marks. Where people had guessed, they also understandably struggled to make any sensible relationship between the data structure and the technique.

**Which two of the data structures in Figure 2 from the previous question would be appropriate for storing volumetric data to be used for Direct Volume Rendering? Explain which would work best for dense data resulting from a geological scan, and which for more sparse data such as a CT scan of an archeological object inside a locked box. In each case, name the data structure and explain its use.**

Answers to this question fell broadly into three categories: people who didn't even try and got 0, people who guessed (and maybe scraped a mark), and people who knew what they were talking about (and got full or nearly full marks).

**Briefly describe a process of Direct Volume Rendering and a process of Indirect Volume Rendering, then compare and contrast the pros and cons of these approaches when used in the context of a medical application such as the detection of a tumour within healthy tissue.**

As above this question was either answered well, or skipped; there didn't appear to be any obvious misconceptions.