

CS 410, Fall 2017, Second Third of Semester Review November 2, 2017

The following is not an exhaustive list of what has been covered, and to be very explicit, the Second Midterm may cover material discussed in class not been mentioned below. That said, hopefully the following will be a helpful as you review the broad topics and techniques we have covered since Midterm 1 in CS410.

1. You should be able to describe the high level and key distinctions the projective pipeline versus ray tracing approaches to rendering 3D scenes.
2. You now have, if you did not already, a basic familiarity with the characterization of visible light in terms of a spectrum where the wavelengths vary between 380 nm to 780 nm.
3. There are several interrelated ideas that you now can explain to a newcomer all related to the single key term "trichromaticity".
4. Consider the following statement: "The color of light is something very different to a physicist or a hyperspectral sensor than it is to a human eye." You should be prepared to give a concrete example of how this statement is true.
5. One of the most fundamental practical concepts regarding color goes by the name "RGB Cube". Fortunately, you now completely understand what this term is describing.
6. As useful as the RGB cube is, it is not always the best way to think about color. An alternative color space is HSV. Having learned what HSV stands and how it relates to the RGB cube you, should be prepared to answer questions that relate to mappings from one space to the other.
7. To exercise your knowledge of RGB color space, you can match up color codes with the 16 named colors established in the HTML 4 specification (Google "named Web safe colors"). In particular, the colors would be: White, Silver, Gray, Black, Red, Maroon, yellow, Olive, Lime, Green, Teal, Blue, Navy, Fuchsia and Purple.
8. Colors are often specified using hexadecimal numbers. For example, #800080 is purple. You should be comfortable answering questions about colors where the colors are specified in this notation.
9. While you certainly should not have memorized the entire catalog of hexadecimal color specifications, you should be comfortable with questions of the style: " Which of these 2 colors is more saturated?" Or alternatively, "Which of these 2 colors appears brighter?"
10. The use of an alpha channel to blend the color of one object with that behind it has now been carefully reviewed and you should be comfortable answering questions about alpha blending.
11. Given its relative simplicity when compared to perspective projection, orthographic projection can be easily overlooked. However, you understand important contexts in which it is most useful and you can define orthographic projection by writing down an example of an orthographic projection matrix.
12. A predator endowed with orthographic vision (orthographic projection) is not fooled by an elephant hiding behind a rabbit. Not so for standard issue predators whose eyesight follow the laws of perspective projection. You should have no trouble sketching an illustration of this admittedly somewhat off-the-wall approach to contrasting orthographic and perspective projection.

13. In the lecture introducing perspective projection, the first perspective projection matrix was derived for the case where the image plane is moved an amount d in front of the focal point. The second formulation placed the image plane at the origin with the focal point behind the image plane by an amount d . The corresponding perspective projection matrices look similar; swapping 1s and 0s between two positions yields one from the other. However, they behave differently and understanding these differences is important.
14. Some people in computer graphics choose to think of perspective projection as a linear operator, while others do not. It matters less who is right than that you can clearly state the crux of each side's argument. (It would be good to ask about this item in our review session since the issue has probably not received the attention it deserves.)
15. The terms "frustum" and "canonical view volume" are used nearly interchangeably when describing the perspective projection pipeline. Both terms now make sense to you, and you should be able to explain them in your own words and with your own simple illustrations.
16. The most useful perspective projection matrix introduced in this course yields a third coordinate that relates to relative distance between the near and far clipping plane. You can recognize this matrix and explain the basics of how it functions. In particular, what is the nature of the mapping between position between the near and far clipping planes and the corresponding value in the third position of 3D points after they are projected into the canonical view-volume.
17. A complete working model of the perspective projection pipeline has been presented using SageMath. Since you have experimented with different camera specifications using this notebook, you can associate how an object – in our case a simple house – appears in the canonical view volume based upon alternative camera specifications.
18. In the perspective projection pipeline SageMath notebook you have a complete symbolic presentation of the transformations used to transform 3D world coordinates into 2D vertices on an image plane. Therefore, you can easily recognize and reason about the significant components in this transformation.
19. You have been given a brief introduction to a low-level operation associated with polygon rendering often described as scan conversion; the process of coloring pixels that belong to a 2D polygon. While our treatment of this topic in CS 410 stops short of presenting a complete algorithm, you are aware first of the critical need to ultimately fill – color – pixels. In addition, you can discuss the most basic terminology and considerations which come into play. These include 'scan-line', 'z-buffer' and finally the principal that no pixel can belong to two polygons.
20. Three kinds of reflectance lie at the heart of computer graphics: ambient, diffuse (Lambertian), and specular. Your knowledge of all 3 is thorough and relates both to mathematical definitions and intuition with respect to identifying which of these are present in a given rendering of a scene. Your intuition is further bolstered by time spent experimenting with the SageMath notebooks which render illuminated scenes containing one or several spheres.
21. Ambient illumination may be thought of in several ways. First, is an extraordinarily convenient and simple hack. Second, a key element in realizing a believable rendering of many scenes. Third, a very coarse approximation to a common condition on the coast of California. (This last part might deserve a brief mention in the review section)
22. Two of the three kinds of reflectance require explicit light sources, and one does not. Be clear about which is which.
23. One and only one of the 3 types are reflections require knowledge about the relationship between the viewer (camera) and the position on the surface whose illumination is being computed. Be clear about which of the three

types of illumination requires this extra information, i.e. direction to the camera.

24. You know how to compute the direction to a light source at $[3, 4, 1, 0]$ specified in homogeneous coordinates. As you think about this light source specification recall two things. First, in Fall 2017 you will not be asked write code using such a light source. Second, answering this question is easier than you might think and if you are in doubt of the correct answer then ask.
25. The computation of diffuse reflection from a point on a surface depends upon the intensity of the light source, the normal to the surface, and the vector indicating the direction from the point on the surface back to the light source. The exact equation and motivating diagram is something you can reproduce from memory.
26. Consider if you would agree with this statement: "The essence of Lambertian reflectance is the way photons are reflected back in the direction of the light source with greater intensity/frequency than they are reflected off to the sides." As you consider your response to the perhaps deceptively simple statement, it would probably be helpful to step through exactly what happens to a photon that hits a surface under the Lambertian model of reflectance.
27. There is one more thing that enters into computing diffuse reflectance, and it is typically expressed using either 3 or 9 constants (K). In terms of the underlying physical model of surface reflectance you can readily explain what these 3 values express, i.e. how they are used. Similarly, you can do the same if 9 values are provided.
28. For a light source located at a position in the world (x, y, z) and a point on a surface (sx, sy, sz) , you can write down the precise values in a vector pointing back to the light source.
29. One might say of diffuse versus specular reflection that one is deep while the other is superficial. While this is admittedly a somewhat odd statement, it captures something of importance that you in turn can now explain.
30. Of the 3 forms of reflection, specular reflection involves the greatest amount of required information. In particular, it requires a vector representing the direction to the light source (also required for diffuse), a surface normal (again same as diffuse), a reflection Ray R that was not required for diffuse illumination, and finally a vector representing the direction to the camera/viewer. Of these, the reflection Ray R is arguably the most involved to compute. Since it is a fundamental element in computing specular reflection, you of course understand the construction of how it is expressed relative to the other known elements.
31. It seems the parameter alpha crops up in many places in computer graphics. It plays a critical role in Phong specular reflectance and is commonly assigned values such as 5 or 50, or even higher. You can explain precisely what alpha is doing both in words and with an equation.
32. There is an essential and very interesting aspect of specular reflection and its associated constant(s) k . Namely, there are 3D modeling packages/languages where k is expressed as a scalar rather than a three-tuple (three scalar constants). Given your understanding of specular reflection, you can now explain why only a single scalar value is a good choice in many circumstances. Likewise, you can explain the flexibility introduced by using a three-tuple.
33. Sometimes in 3-D modeling for computer graphics being "two-faced" is not such a bad thing. Be ready to explain what such an off-handed remark actually means.
34. Two concepts are easily confused. The first is the field of view of a pinhole camera. The second is the resolution, or pixel density, of an image. In a well-built rendering system, the controls used to modify the field of view are clearly separate from the controls used to change the pixel density. Having now built your own rendering system, you're familiar with this distinction, and therefore should have an easy time explaining it to others and fielding questions relating these two concepts.

35. An entire lecture was given over to illustrating basic aspects of 3-D modeling using the program SketchUp. One skill you should now have taken away from that lecture is relating how a 3D model (a simple model) appears visually and appears in a standard file format such as the .OBJ format.
36. It would be a lot to expect anyone to keep the detailed procedures for many of the operations supported in SketchUp in their heads. However, a good understanding of the most common and essential operations supported by such a tool represents basic knowledge anyone working in computer graphics ought to have. Therefore, you are comfortable answering questions regarding the basic functionality demonstrated in lecture.
37. Given the direction of a ray defined by a vector arriving at a point, as well as the 3D coordinates of the point, and finally the surface normal, you are able to choose between a series of alternative possible reflection rays provided the choices are qualitatively distinct. In other words, incorrect choices may be ruled out based on some clear conflict such as pointing in the overall wrong direction.
38. Both diffuse and specular reflection have dot products as key to their computation. In both cases, a negative value for the dot product has important geometric implications and also requires explicit handling in code. As you go implement a ray tracer, and when you go to explain the process to others, you are comfortable now with the meaning and proper handling of those situations where the dot product yields a negative number.
39. For recursive ray tracing two distinct concepts have been introduced in general terms as well as precisely in SageMath code. These two concepts are attenuation and reflectivity. Having experimented with both in the SageMath notebook on recursive ray tracing discussed in lecture on October 31st you are comfortable with how each operates and how they are distinct from each other.
40. With all other camera and scene parameters left constant, you are now comfortable explaining what happens when recursion depth is increased. In particular, you can correctly argue the following proposition: with increased recursion the intensity of any given pixel must remain constant or increase.