CS 410, Fall 2013, Things You Should Know going into Midterm 1
September 26, 2013

Disclaimer: this study guide is provided in the spirit of assisting students to organize and recall what has been covered. It is not a contract and material not appearing here but covered in lecture or assigned portions of the textbook may appear on the exam.

1. You should be able to describe the high level and key distinctions between the 2 approaches to rendering images in graphics: specifically the projective pipeline versus ray tracing.

2. The mathematics of 3-D modeling, both of objects and cameras, rests squarely upon a solid working understanding of the following 3 terms: scalar, vector and point. Your understanding of computer graphics at this point means you're able to succinctly define each of these 3 mathematical concepts.

3. The concept of a vector space is straightforward, for the sake of this class you may define it simply as the space of all vectors which may be defined using a fixed number of scalars: typically 3. The concept of an affine space is somewhat more difficult to hang onto, but you now know the one thing that affine spaces have that the vector spaces do not, and you can use it to remind yourself about the difference between these 2 concepts.

4. You now appreciate that the dot product has significance for geometric modeling at many levels. At the most pragmatic level, you certainly know how to compute the dot product between 2 vectors. At a much deeper level, the dot product represents the transition from affine spaces to Euclidean spaces. You're now in a position to explain why.

5. You have now seen how basis vectors play a truly fundamental role in how we express the geometry of object models and cameras. Here's a thought question you should be ready to tackle. In order to span a space, is it absolutely essential that basis vectors be orthogonal? Another way to think about this question is as follows, if it is not a requirement that basis vectors be orthogonal, why does it appear that we place such a high premium on selecting orthogonal basis vectors in practice.

6. The simple English phrase "the same point" is more ambiguous and perhaps fraught with peril than you would have initially imagined before taking this course. At a minimum, there are 2 fundamentally distinct and separate interpretations of the phrase. You should now be prepared to describe each of these and contrast them with each other.

7. You are now aware that there are 2 rather distinct ways of introducing the concept of rotation in the 2-D plane. Most references (e.g. Wikipedia) simply state the rotation matrix in terms of the cosine and sine of an angle theta. There is nothing incorrect about this approach, but it does not really help one understand what is actually taking place when a 2-D point is rotated. Fortunately, you can now take your knowledge of the dot product and derive a 2-D rotation matrix without ever having to touch a cosine or a sine.

8. There are many ways to motivate homogeneous coordinates, fortunately you are now in a position to quickly provide a pragmatic motivation in terms of chaining together, or perhaps one should say composing, sequences of transformations. Further, if making this argument to another it is best to illustrate using concrete examples of commonly named types of 2-D transformations.

9. It goes without saying that at this point in the course you can write down examples of 2-D rotation, translation, scaling, non-uniform scaling, and sheer matrices - all in homogeneous coordinates. Just as

important, if you are presented with a set of 2-D points drawn before and after the application of one or several of these transformations, you can now work out by looking at the relative positions of the points pre- and post-transformation which operation or combinations of operations were carried out.

10. Perhaps one odder aspect of learning computer graphics is the need first to learn about Euler angles and subsequently to learn why in practice they are virtually never the best way of capturing the rotation of an object or a camera. Fortunately, you can both explain what Euler angles are and why they are not heavily used.

11. In general, one technique for determining a 2-D transformation is to view the terms in the transformation matrix as unknowns and the coordinates of points before and after the transformation as constraints. For arbitrary combinations of rotation, translation, and scaling, this is not typically the sort of thing that would be asked in in an exam. However, for simpler types of transformations, perhaps a combination of translation and scaling, it is entirely reasonable exercise.

12. If you are asked over lunch in a job interview to compute the cross product between the vectors $(1, 2)$ and $(2, 3)$, be certain that you would know why your reaction should be one of puzzled amusement, aware of the fact that you are being baited into possibly revealing a basic misunderstanding of geometry.

13. There are 2 ways you should have the cross product committed to memory. The 1st is to be able to explain in English the geometric interpretation of the resultant vector. The 2nd is the algebraic approach in which you can actually compute the cross product of 2 vectors.

14. If one were to say there's something special about the cross product between 2 unit length and perpendicular vectors, you fortunately could quickly what is special about this case.

15. In 3 dimensions as opposed to 2 dimensions, it is a bit harder to make up matrices that represent valid 3-D rotations. That said, is not all that hard, and you now know how. You should be able to explain by example the rules governing the construction of a valid rotation matrix.

16. Just as computer graphics folks don't typically think much of Euler angles, many of us have a fondness for the axis angle approach to specifying 3-D rotations. Fortunately, you can now illustrate this by showing how to rotate $45°$ about the vector $(2, 2, 1)$.

17. Given its relative simplicity when compared to perspective projection, orthographic projection can be easily overlooked. However, you do understand its uses, and you could certainly write down an example of an orthographic projection matrix.

18. Our lectures on modeling cameras have used phrases such as "bear in the box" and "pinhole room". What underlying critical concept rests beneath the illustrations that have in turn given rise to these phrases?

19. A predator with orthographic sight (orthographic projection) would not be 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.

20. In the lecture introducing perspective projection, the 1st perspective projection matrix was derived for the case where the image plane is moved an amount d in front of the focal point (perspective reference point). This was then contrasted with an alternative formulation involving an image plane at the origin, the perspective reference point behind the image plane by an amount d, and finally points in the world being imaged in front of the image plane an amount z. Superficially, the 2 matrices look extremely similar, one can

even arrive at 1 starting from the other by simply swapping ones and zeros between two positions within the matrix. That said, they behave very differently, and fortunately you're in a position now to explain the difference.

21. Some computer graphics folks choose to think of perspective projection as a linear operator, while others do not. It matters less who is right then that you can clearly state the crux of each side in this argument.

22. The 2 terms "frustum" and "canonical view volume" are used nearly interchangeably when describing 3-D clipping. Beth terms now make sense to you, and you should be able to explain them in your own words and with your own simple drawings. Doing so includes the critical concepts of near and far clipping planes.

23. An infinite line in 2-D may be described as all points a fixed distance from the origin measured in a particular direction. You can now move back and forth between this English description and mathematical formulation which it implies.

24. On the way to understanding how a ray is parameterized and represented, one typically has to first master the concept of a parametric line. Indeed, you are now able shift between three common line/line-segment representations: slope-intersept form, implicit form, and parametric form.

25. The meaning of the parameter in a parametric line segment can take on one of two possible interpretations depending how the form is derived. The distinction is fundamental and one you can explain to others and answer questions about in a job interview.

26. In Ray tracing, how is the direction of the ray determined? When explaining this to someone it is necessary to be precise.

27. There is a minor but important difference between 2 alternative ways of representing a ray. In one, less work is required by the code when moving from pixel to pixel to pixel. In the other, the t-values associated with the first intersection may be interpreted directly as distances into the world. Within this context, you should be able to precisely describe the 2 representations and highlight why the 2 behave differently.

28. 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. You now understand the relationship between the three essential compenents of making objects appear larger.

29. There is a kind of word problem you can now easily solve: a space ship leaves earth in a direction U, how far has it traveled when it passes closest to a planet at position P. Assume the earth is the origin.

30. There is a clever way of determining the intersection of a sphere and a ray. Since you have already implemented this as part of your programming assignments, you're aware that solving a general problem without the aid of a calculator would be annoying. That said, the construction of the algorithm is elegant and important, and you certainly could work out a concrete numerical example if the constituent components of the example were well enough chosen to make ensuing arithmetic straightforward.