You must answer all questions.

For full credit, an answer must be both correct and well-presented (clear and concise).

If you feel a question is ambiguous, state any assumptions that you need to make. Also, several of the questions are "essay" questions. For these, there are many correct answers. It is more important that you provide a good argument for the answers you give, than that you give the "most correct" answer. Sometimes, we are particularly looking for your ability to make a clear and concise argument based on things you are aware of, rather than to see if you can find the best possible answer, or have seen all possible research on the topic.

Problem 1: Rotations and transformations (3 questions)

1. We want to construct a rotation that transforms the unit vector along the x-axis $e_1=(1,0,0)$ into an arbitrary unit vector v=(x,y,z). Construct the quaternion that corresponds to this rotation – of course, your answer will be a function of the target coordinates (x,y,z).
2. Given the 3x3 matrix of a purely rotational transformation, describe how we can determine the axis of rotation that it corresponds to.
(You may not assume that the quaternion form of the transformation is readily available).
3. A linear (more accurately, affine) transformation is encoded in the mapping T[x] := **M**x+**t,** where **M** is a 3x3 matrix and **t** is a 3D vector.
    a. What would it mean, intuitively, if all 3 eigenvalues of **M** were real numbers?
    b. In addition to (a), what would be the significance of the corresponding eigenvectors being mutually orthogonal?
    c. How can we check if such a transformation has the property that the Euclidean length of any vector remains unchanged after the transformation has been applied?

Problem 2: Geometry representation and manipulation (2 questions)

1. Implicit surfaces represent the surface of an object as the zero isocontour of a scalar field. Signed distance fields (also referred to as level sets) additionally have the property that the absolute value of the scalar field measures the distance to the object surface. Imagine having two similar, but not identical models, each encoded as an implicit surface. We want to morph between the two, in a smooth fashion, and describe the result as an implicit surface as well. In which instances could we expect simple linear interpolation of the respective scalar fields to be a good solution (or approximation)? Can you describe a scenario where simple linear interpolation of the scalar fields would risk artifacts, or otherwise unacceptable results?

2. You are developing computer modeling software that allows you to design and style digital models of women's dresses, for a clothing company. The design department is providing you with surface models of dresses, corresponding to early prototypes, that they authored using third party software. Their triangulated surface models are given to you simply as a list of vertex triples $(i_1, j_1, k_1), (i_2, j_2, k_2), \ldots (i_n, j_n, k_n)$ containing the integer vertex indices of each triangle and, additionally, and a list containing the 3D coordinates of all vertices.

   [Note: the specifics of the file format encoding the information above are not significant. What is significant is that you are not <u>*explicitly*</u> *given any additional information such as pairs of adjacent triangles, incident triangles on a vertex, etc. If you need any such information, you should explain how you compute it.]*

   a. Prior to inputting this mesh into your pipeline, you want to perform some rudimentary sanity checks that the input mesh is of "good" quality. Among those, you want to ensure that the garment has exactly 4 holes/openings, one at the bottom, one at the neckline, and one at each sleeve. How would you implement that check using only the information provided in the described input format?

   b. Suggest additional sanity checks that you might include, to help ensure the garment model is ready for follow-up modeling, texturing, and possible simulation.

<u>**Problem 3: Diffusion curves (the "modern paper")** (3 questions)</u>

You were asked to read the paper:

"**Diffusion Curves: A Vector Representation for Smooth-Shaded Images**"
by Alexandrina Orzan, Adrien Bousseau, Holger Winnemöller, Pascal Barla, Joëlle Thollot, and David Salesin, *ACM Transactions on Graphics* 27, no. 3 (Proceedings of SIGGRAPH 2008)

Note: there have been a number of follow-up papers that may address some of the issues related to the questions asked below. If you have read about something you think is an answer to one of the questions, you can describe the solution from a paper (please refer to the paper). However, these questions do not expect you to have read those papers. They can be answered based solely on an understanding of the paper, and the expected background material in the qual. In fact, for some of the questions, the newer methods might be an overkill (the problems are simpler than those discussed in the newer paper).

1. In the paper, the curves are defined as cubic Bezier curves, and colors are placed at specific parameter values (t) and linearly interpolated between those. Imagine that we only allow placing colors at the <u>interpolating control points</u> (e.g. the first and last control points of each Bezier segment).

   a. There are a number of different ways we could interpolate the colors. For example we could (i) employ linear interpolation using the curve parameter, (ii) we could use arc length interpolation, or (iii) use Bezier interpolation after doubling the first and last point of each curve. Compare and contrast these approaches. Give concrete examples of where they would produce different results. (or argue that they give the same results)

b. To get around the limitation of only being able to place colors at control points you could imagine "cutting" the curve into two segments (this now yields 3 control points, since the new positional endpoint would be shared among the two new segments). For each of the three different interpolation types, say whether or not the new curve (consisting of 2 segments) would exactly reproduce the results (positions and colors) of the old curve. Be sure to give your reasoning.

2. The paper discusses painting in 2D by drawing curves. Suppose that you wanted to paint on a 3D triangle mesh. You can assume that the "curve" is a "polyline" (a sequence of mesh vertices), and that it is sufficient to merely determine colors for the vertices of the mesh.

   Discuss the issues in creating the new method for diffusion curve painting on a 3D triangle mesh. Question (3) below may give you some hints. However, you should have at least some issues that will be true even for good quality meshes.

3. Some of the issues in making diffusion curves work for meshes come from the desire to get decent results with a poor quality mesh. Describe some issues relating to making diffusion curves work on poor quality meshes. Discuss how various methods for improving mesh quality might address these.

Problem 4: Mesh-based deformers (4 questions)

A recent approach to character animation is to enclose the geometry of the character (typically a polygonal mesh) within a simpler polyhedron, called a cage. As the cage is manipulated, the geometry is deformed accordingly.

1. What are the advantages and disadvantages of having animators manipulate the character using the cage rather than the surface mesh itself? (Hint: there are good and bad traits in both approaches)

2. One way to implement the cage-based deformations is to divide the control polyhedron into tetrahedra, and use barycentric interpolation to position each vertex of the mesh within the tetrahedron that contains it. A second way to implement cage-based deformations is to use some form of generalized Barycentric coordinates to define a single coordinate system over the entire control polyhedron. Discuss the advantages and disadvantages of each approach. (Hint: there are good and bad traits in both approaches)

3. Harmonic coordinates are one popular technique used to create the generalized barycentric coordinates for skinning. What nice properties of Harmonic coordinates make them particularly useful for skinning?

4. Common skinning techniques use an internal skeleton, and perform a weighted average of the different coordinate systems (associated with individual bones) for each vertex. What are the pros and cons of cage-based control compared to this more common skeletal-based skinning?

   [Note: Often skeletal skinning uses linear blending (or weighted averaging), which causes bad artifacts. If you list these kinds of visual artifacts as a con (and therefore a pro of cage-based skinning), you must explain why it is easier to address these issues in a cage-based approach than a skeleton-driven approach. Remember that many of the artifacts have been addressed by more recent skeleton-driven approaches such as Dual Quaternion skinning.]

Problem 5: Color theory (4 questions)

You should know that humans with "normal" color vision are tri-chromats. Color-vision deficiency (CVD) means that people are only bi-chromats. A very rare genetic condition causes some people to be tetra-chromats (meaning they have 4 different types of cones).

In tri-chromatic vision, the cones types are L (long), M (medium) and S (short). These roughly correspond to red, green, blue sensitivity. Each cone's response is a bump, with its center in a different place. Each type of cone responds to all frequencies of light, except they respond more strongly to colors near their centers.

Real tetrachromats are rare enough that we don't know much about the response of the fourth cone.

1. Even for a tri-chromat, a standard 3-color (RGB) display may not be able to create all possible color responses. In fact, any physically realizable 3-channel display cannot. Explain why.
2. Some manufacturers have suggested making 4-color displays to improve color fidelity. There have been proposals to add a Yellow (between red and green in the spectrum) (e.g. RYGB) and one to add Violet (even higher in frequency than blue RGBV). Explain how each of these displays might help provide a more complete color experience. Describe the pros and cons of the two approaches. For simplicity, you may assume that each of the 4 light sources (or 3 light sources in the standard monitor) are lasers that produce a single frequency of light.
3. Suppose you had to build a display for a tetra-chromat, but you didn't know the response of their fourth cone type (but you did know the first three cones were the same as in normal vision). You have to add a fourth color to RGB, but you need to make a choice of that color that is not informed by the color sensitivity of the fourth cone type. Would this monitor help the tetra-chromat? (assuming you picked a color that was not RGB, but wasn't necessarily the color of the 4th cone)
4. Two colors appear gray and isoluminant (equal brightness) to a tetrachromat. Is it possible for them to appear to have different brightness values to a tri-chromat? Explain why, or why not.

Problem 6: Rendering Systems and Hardware (1 question)

The original Reyes renderer (which became the Renderman) was an important high-quality rendering system, which is still the basis of many modern, high-end software renderers. However, the basic ideas that it introduced were very influential on future hardware rendering systems as well. Give some examples of how ideas from Reyes/Renderman are manifested in modern hardware graphics systems.