**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.**

## Question 1: The Fall Forest (3 sub-questions)

Imagine a Wisconsin forest in the autumn at the height of the "fall colors." The leaves have changed to the wonderful bright colors of the season. The sunlight filters through the leaves (not only are leaves not completely opaque, but light bounces around in the trees) giving a very special quality to the light in the forest. When the leaves blow, and the leaves move, the light is "animated" and can be particularly beautiful. In this question, we consider trying to recreate the wonderful experience of the Autumn forest.

You may assume that there are methods that allow you to generate the geometry of the trees. (good methods exist, although they are not on the qual reading list). You may assume that the generator provides you with the detailed shapes and positions of the leaves and branches.

Essay questions:

1A. What makes this scene a challenging one to render effectively (both in terms of the quality of appearance, as well as performance). Give some examples of existing techniques and why this scene might challenge them.

1B. What are the issues in simulating the motion? (both in terms of the quality of movement, as well as performance). Give some examples of existing techniques and why this scene might challenge them.

1C. Choose one of the set of challenges (1 or 2). Discuss how you might approach it by extending current techniques. Discuss the available approaches and discuss their pros and cons. If you were to try a research project to do this simulation, what would you choose and why?

## Question 2: Re-Sampling Discrete Structures

The combination of low-pass pyramids and multi-way interpolation is a common solution for sampling discretely represented images (technically, this is re-sampling). For example, for surface texturing, MipMaps (a specific type of low-pass pyramid) combined with tri-linear interpolation is the de facto solution. Such solutions are desirable because they provide efficiency through pre-computation and quality tradeoffs.

For 2A and 2B, assume that the original image data is regularly sampled on a grid.

2A: For the specific case of mip-mapped based, tri-linear interpolation, explain situations where these tradeoffs may not be acceptable. Give situations for both the performance assumptions failing to be valid, as well as where the quality assumptions may not be valid. For each, please explain the technical problem, but also give an example of a realistic application for which this may be a problem.

2B: Consider the application of the approach to higher dimensional data (e.g. 3D volumes). How do these approaches extend? What new issues and tradeoffs arise as the dimension of the data goes up?

2C: Consider what happens when the input data is not regularly sampled on the grid. Can you suggest ways of applying the pre-filter-pyramid plus interpolation approach to these cases? You can restrict your thinking to 2D if that is helpful.

## Question 3: Discrete Surface Models

Several established techniques exist for encoding  the geometry of surface models, each with its individual advantages and drawbacks. Possible options include:

1. Tessellated surface models, e.g. surface meshes with triangular, quadrilateral or polygonal elements.
2. Spline surfaces.
3. Subdivision surfaces.
4. Implicit surface representations (for example, defining a surface as the set of points (x,y,z) in the 3D space that satisfy an equation f(x,y,z)=0)
5. Constructive Solid Geometry (CSG) [Note: Although CSG is normally associated with descriptions of volumetric objects, we are simply using it in this context as a method for representing the *boundary surface* of such an object]

For ***each*** of these five representations, describe one scenario where the use of the technique would be advantageous (and explain why), ***and*** one scenario where this approach should be avoided (and explain why not). A non-exhaustive list of scenarios that you may choose to consider is:

- Architectural models, such as buildings, interior spaces and engineering structures.
- Items of clothing.
- Objects that are intended to be viewed at various distances and/or levels of magnification
- Discrete representations of models that are acquired via (3D) scanning hardware.
- Surfaces of fluids, that may potentially be animated over time.

## Question 4: Rendering and randomness

A graduate student is planning to write a renderer and intends to make it as photo-realistic as possible. Her goal is to produce images that depict as many of the lighting effects seen in the real world as possible.

Her advisor (who might not be as knowledgeable about best practices in graphics and rendering) insists that she write her renderer without using any random number generator, as he believes this will mean that her results are more deterministic and reproducible. This precludes her use of several modeling and rendering algorithms.

Describe the consequences that this design decision may have on the renderer that the student develops. Consider aspects such as image quality, impact on computational time, limitations on the effects that can be reproduced, or risks of visible artifacts under certain scenarios.

[Note: Although your reading list does not exhaustively discuss all algorithmic aspects of rendering techniques that might be related to randomness (or lack thereof) you are asked to use critical thought to answer this question to the best of your knowledge]

## Question 5: Simulation and Modeling

A. The introduction of implicit time integration schemes for simulation in Computer Graphics applications has enabled significant advances in terms of tractable resolution, robustness, and interactivity. Nevertheless, implicit methods and their ability to handle large time steps are not a panacea. Describe 3 reasons why an implementer may choose to either (a) use an explicit time integration scheme instead, or (b) be quite conservative with the time steps taken, even if his implicit scheme would allow him to be more aggressive.

B. We want to model (and animate) the process of a surface being torn or fractured. Examples would be a glass pane being shattered, a piece of paper being shredded or a piece of cloth being torn. What would be some reasonable options for representing the geometry of such a dynamically changing surface? Include considerations such as rendering, continuity of animation and versatility to encode different tear patterns in your answer. [Note: Although you are free to consider how a certain representation might facilitate physics-based simulation of tearing/fracturing, you are not necessarily restricted to such a scenario. For example, you may ask yourself what geometric representation might be appropriate for a tearing animation that is procedurally generated or artistically keyframed by an animator, or for a real-time game that includes glass shattering effects].