Final Exam Review

15-462 / 15-662 Computer Graphics

Lecture 1: Introduction

- For any given setup where we place a camera in the environment, pointing down any of the main coordinate axes (x, y, or z), compute a projection of points in the world onto an image plane.
- Write an algorithm for drawing lines that handles all edge cases (i.e., including edges that are exactly horizontal or vertical).

Lecture 2: Linear Algebra Preview (Part 1 of 2)

- **What information does a vector encode?**
- **Draw a geometric representation of each rule that vectors seem to obey.**
- **What is a Vector Space? Give an example of a Vector Space.**
- **Can a function be a vector? Explain**
- **Add and scale vectors.**
- **Add and scale functions**
- **What is the norm of a vector?**
- **Associate each property of a norm with a geometric interpretation.**
- **Compute the Euclidean norm in Cartesian coordinates.**
- **Compute the L2 norm of a function.**
- **Associate each property of the inner product with a geometric interpretation.**
- **Compute the inner product in Cartesian coordinates.**
- **Compute the inner product of two functions.**

Lecture 2: Linear Algebra Preview (Part 2 of 2)

- Give properties and an example of a linear map.
- Define span and basis.
- Compute an orthonormal basis from a set of vectors.
- Be able to use Gram-Schmidt orthonormalization.
- Know that orthonormalization of functions can be done by decomposing them into sinusoids.
- Be able to solve a simple system of linear equations, depict it geometrically, and represent it in matrix form.
- Be able to represent a linear map in matrix form.

Lecture 3: Vector Calculus Preview (Part 1 of 2)

- Euclidean norm is any notion of length preserved by rotations/translations/reflections of space. Be able to calculate it for a vector of any dimension.
- Compute the inner product of two n-dimensional vectors. What is the geometric meaning if an orthonormal basis is used?
- Compute the cross product of two three-dimensional vectors. What is the geometric meaning of this cross product? What is the geometric meaning of its magnitude?
- Use the cross product to do a quarter rotation of a vector within a plane.
- Represent the dot product using matrix notation.
- Represent the cross product using matrix notation.
- Understand what a determinant measures. What does the determinant of a linear map tell us? Give an example.
- What is a directional derivative?
- Compute the gradient of a function.

Lecture 3: Vector Calculus Preview (Part 1 of

- Understand gradient as the best linear approximation and the direction of steepest ascent.
- When is the gradient not defined?
- Be able to express gradients of simple matrix expressions (Slide 31).
- What is a vector field? Give an example.
- Be able to compute divergence, curl, and the Laplacian of a vector field. Also be able to express the meaning of these terms geometrically, for example by drawing a diagram.
- What is the Hessian? Be able to compute the Hessian of a function.

Lecture 4: Drawing a Triangle

1. How should we choose the correct color for a pixel? There is not an exact right answer. However, you should be able to discuss some of the issues involved.
2. What is aliasing, and what artifacts does it produce in our images and our animations?
3. One form of aliasing is where high frequencies masquerade as low frequencies. Give an example of this phenomenon.
4. Suppose we have a single red triangle displayed against a blue background. Does this scene contain high frequencies?
5. What does the Nyqvist-Shannon theorem tell us about how image frequencies relate to required sampling rate?
6. The practical solution on your graphics card for reducing aliasing (i.e., for antialiasing) is to take multiple samples per pixel and average to get pixel color. Try to use what we learned about sampling theory to explain as precisely as you can why taking multiple samples per pixel can reduce aliasing artifacts.
7. Be able to write an implicit representation of an edge given two points.
8. Be able to use the implicit edge representation to determine if a point is inside a triangle.

Lecture 5: Transformations (Part 1 of 2)

1. Which of the following operations are linear transforms: scale, rotation, shear, translation, reflection, rotation about a point that is not the origin?
2. Express scale as a linear transform
3. Express rotation as a linear transform
4. Express shear as a linear transform
5. Express reflection as a linear transform
6. Express translation as an affine transform
7. Explain how to compute rotation about a point that is not the origin
8. Know what makes a transform linear vs. affine
9. Know how to build transformation matrices from start and end configurations of your object

Lecture 5: Transformations (Part 2 of 2)

- Create 2D and 3D transformation matrices to perform specific scale, shear, rotation, reflection, and translation operations
- Compose transformations to achieve compound effects
- Rotate an object about a fixed point
- Rotate an object about a given axis
- Create an orthonormal basis given a single vector
- Understand the equivalence of [x y 1] and [wx wy w] vectors
- Explain/illustrate how translations in 2D (x, y) are a shear operation in the homogeneous coordinate space (x, y, w)

Lecture 6: 3D Transformations

1. Create a rotation matrix to rotate any coordinate frame to xyz
2. Create the rotation matrix to rotate the xyz coordinate frame to any other frame
3. Know basic facts about rotation matrices / how to recognize a rotation matrix
   A. Rows (also columns) are unit vectors
   B. Rows (also columns) are orthogonal to one another
   C. If our rows (or columns) are u, v, and w, then uXv=w
   D. The inverse of a rotation matrix is its transpose
4. What is the problem of gimbal lock? Give an example where this problem occurs.
5. How does using quaternions solve this problem?
6. Know that every rotation can be expressed as rotation by some angle about some axis.
7. Know how to go between quaternions and axis-angle format for rotations.
8. Know how quaternions derive from higher dimensional complex numbers.
9. Be able to work out quaternion multiplication from the complex number representation of a quaternion.

Lecture 7: Projection and Textures (Part 1 of 2)

- Create a projection matrix that projects all points onto an image plane at z=1
- Propose a projection matrix that maintains some depth information
- Understand the motivation behind the projection matrix that projects the view frustum to a unit cube
- Be able to draw / discuss the details of the view frustum
- Prove that a standard projection matrix preserves some information about depth
- Interpolate colors using barycentric coordinates
- Compute barycentric coordinates of a point using implicit edge functions
- Compute barycentric coordinates of a point using triangle areas
- Estimate the location of a point inside a triangle given its barycentric coordinates
- Estimate the location of a point outside a triangle given its barycentric coordinates
- Estimate barycentric coordinates of a point from a drawing.
- Show that interpolation in 3D space followed by projection can give a different result from projection followed by interpolation in screen space. In other words, explain why interpolation using barycentric coordinates in screen space may give a result that is incorrect.
- How, then, can we obtain a correct result using interpolation in screen space?

Lecture 7: Projection and Textures (Part 1 of 2)

- Textures are used for many things, beyond pasting images onto object surfaces.
  - Normal maps (create appearance of bumpy object on smooth surface by giving false normal to the lighting equations)
  - Displacement maps (encode offsets in the geometry of a surface, which is difficult to handle in a standard graphics pipeline)
  - Environment maps (store light information in all directions in a scene)
  - Ambient occlusion map (store exposure of geometry to ambient light for better representation of surface appearance with simple lighting models)
  - Can you think of / discover others?
- Know how to interpolate texture coordinates
- Know how to index into a texture and compute a correct color using bilinear interpolation
- Be able to create a mipmap and store it in memory
- Be able to compute color from multiple levels of mipmaps using trilinear interpolation
- What is the logic behind selecting an appropriate level in a mipmap?
- What can happen if we select a level that is too high resolution?   too low resolution?

Lecture 8: The Rasterization Pipeline

- What is the depth buffer (Z-buffer) and how is it used for hidden surface removal?
- Where does the depth for each sample / fragment come from? Where is it computed in the graphics pipeline?
- Is the depth represented in the depth buffer the actual distance from the camera? If not, what is it?
- What is the meaning of the alpha parameter in the [R G B a] color representation?
- Be able to use alpha to do compositing with the "Over" operator.
- Is "Over" commutative? If not, create a counterexample.

$$C = \alpha_B B + (1 - \alpha_B)\alpha_A A$$

- What is premultiplied alpha, and how does it work?
- Be able to use premultiplied alpha for "Over" composition.
- Why is premultiplied alpha better?
- How do we properly render a scene with mixed opaque and semi-transparent triangles? What is the rendering order we should use? When is the depth buffer updated?
- Draw a rough sketch of the graphics pipeline. Think about transforming triangles into camera space, doing perspective projection, clipping, transforming to screen coordinates, computing colors for samples, computing colors for pixels, the depth test, updating color and depth buffers.

Lecture 9: Introduction to Geometry

- List some types of implicit surface representations
- What types of operations are easy with implicit surface representations?
- List some types of explicit surface representations
- What types of operations are easy with explicit surface representations?
- What is CSG (constructive solid geometry)? Give some examples of CSG operations.
- What type of representation is best for CSG operations?
- Describe how to do union, intersection, and subtraction of geometry using simple operators on a surface representation.
- What is a level set representation? Is it implicit or explicit? (Why?) When is it useful?
- What types of splines are common in computer graphics?
- Why are they popular? What properties make them most useful?
- Derive the equation on the slide labeled "Bézier Curves — tangent continuity" from the definitions given on the previous slides. Draw a diagram to illustrate any terms you use.

Lecture 10: Meshes and Manifolds

- What is a manifold surface?
- Distinguish manifold from non-manifold surfaces
- What does this statement mean: "A manifold polygon mesh has fans, not fins"? Make sketches to illustrate your answer.
- Can a manifold surface have a boundary? Give an example.
- Describe how to store mesh information in vertex, edge, and face tables. Give an example. What are good and bad points of this data structure?
- How would you store a mesh using incidence matrices? What are good and bad points of this data structure?
- What do you need to store in a halfedge data structure? What are good and bad points of this data structure?
- How can you find all vertices in a face with the halfedge data structure?
- How can you find all faces that contain a vertex with the halfedge data structure?
- BONUS: Think of an algorithm to traverse every face in a manifold using the halfedge data structure.

Lecture 11: Geometry Processing

- List practical applications that you can relate to that require good geometry processing algorithms.
- Give criteria for what makes a good quality mesh. Be sure to state your assumptions (e.g., good quality for what purpose?)
- Give pseudocode for one iteration of Catmull-Clark subdivision
- Give pseudocode for one iteration of Loop subdivision
- Explain the important properties of various subdivision algorithms (interpolation, continuity, behavior at the boundaries)
- Be prepared to calculate vertex updates in a simple example of Loop or other subdivision. (The vertex weighting masks will be given to you.)
- Understand how the K matrix of the Quadric Error error metric encodes squared distance to a plane. How can it encode the sum of squared distances to many planes? How is this idea used in generating a good error metric for mesh decimation using edge collapse?
- Describe some techniques for improving the quality of a mesh to make it more uniform and regular.

Lecture 12: Geometric Queries

- Express distance from a plane, given a point on the plane and a normal vector
- Find the closest point on a plane
- Compute ray-triangle intersection, including checking whether the ray passed through the inside of the triangle.
- Be prepared to compute ray-primitive intersection for other primitives (e.g., a sphere) and primitive-primitive intersections (e.g., triangle-triangle or line-line)

Lecture 13: Accelerating Geometric Queries

- Explain the differences between primitive partitioning vs. space partitioning for acceleration data structures. Give examples of each. Give pros and cons for each.
- How would you update a bounding volume hierarchy when object move? How would you update an octree?
- Know how to construct a bounding volume hierarchy (bounding spheres or bounding boxes), a quadtree (2D) or octree (3D), and a KD-tree.
- Know how to traverse any of these trees to test for ray-object intersections along a ray (e.g., to find the first object hit by the ray or to test visibility).

("Lecture 14" was just the Midterm Review)

Lecture 15: Color

- How would you describe the emission spectrum of a light source?
- What are the rods and cones? Which are present in greater number in the human eye? What can you say about how they are distributed in the retina?
- Since color is best described by a spectrum of emissions, why is it that we can get away with just three values for color (e.g., R-G-B)?
- What are additive and subtractive models for color and when are they used?
- Describe some of the different color spaces that are used to express color.
- An alternative to RGB color space is the CIE color space, with X, Y, and Z primaries. What is Y in this color space? What problem with RGB color space does the CIE color space solve?
- Given a color space expressed by some three-dimensional basis it be converted into any other basis through a linear operation (True or False)
- What is gamma correction? Give an example where gamma correction is useful.

Lecture 16: Radiometry

- Visible light consists of a small range of wavelengths along the spectrum from gamma rays to radio waves.
- Energy of a photon depends on wavelength (and speed of light, and Planck's constant)
- What is radiant energy? ..flux? .. irradiance?
- Why (how) does irradiance depend on the angle between the light source and a patch of surface area?
- How does irradiance fall off with distance from the light source?
- What is a solid angle?
- What is radiance, and how many dimensions do you need to capture the radiance in a scene (i.e., to capture a light field)?
- What effect is an ambient occlusion map trying to capture?
- What is radiant intensity?
- (Bonus) Characterize the spectral signatures of different familiar light sources.
- (Bonus) Figure out how to read a Goniometric diagram

Lecture 17: The Rendering Equation (Part 1 of 2)

- **Ray tracing algorithms make use of radiance estimates to build up an image of a scene. What is radiance? Why is radiance the fundamental quantity of interest? How are radiance estimates used to compute the color that would be observe by a camera at a point on a material surface?**
- **What is the difference between radiance and irradiance? Between incident and exitant radiance?**
- **Explain and illustrate with a sketch all of the terms in the Rendering Equation.**
- **Draw diagrams to illustrate reflection for (1) a perfectly specular (mirror) surface, (2) a glossy surface, (3) a diffuse surface, (4) a retroreflective surface**
- **What is a BRDF? What are the inputs and outputs of a BRDF function?**
- **How would you measure a BRDF? If you were to mount a camera and light source on two robot arms, how many degrees of freedom (joints) would you need to do these measurements?**

Lecture 17: The Rendering Equation (Part 2 of 2)

- Given a ray and a surface normal, calculate the direction of perfect reflection
- Given a ray, a surface normal, and indices of refraction, calculate the direction of perfect transmission using Snell's Law.
- What is total internal reflection? Give a detailed example of how / when it can occur.
- What is Fresnel reflection? Sketch curves to illustrate the effect for dielectrics vs. conductors as we have seen in class. Label your axes. Informally, what does this effect show?
- What is subsurface scattering?
- How can we extend the idea of BRDF to subsurface scattering? What additional parameters must be sampled?

Lecture 18: Numerical Integration

- In rendering (global illumination), what are we integrating, i.e., what integral do we want to evaluate?
- How do we use the trapezoid rule to integrate a function?
- How does work increase with dimensionality of our function?
  - This is why we typically use Monte Carlo integration in graphics!
- Give a high level overview of the process of Monte Carlo integration
- What is a probability density function (PDF)?
- What is a Cumulative Distribution Function (CDF)?
- The Inversion Method can be used to correctly draw a sample from a PDF.
  - Sketch the overall step by step process for using the Inversion Method
- What is rejection sampling? Show how to use rejection sampling to sample area of a circle, volume of a sphere, directions on a sphere, and solid angles from a hemisphere.

Lecture 19: Monte Carlo Ray Tracing

- What is Expected Value? .. Variance? .. Bias?
- What is Importance Sampling? How is it used in cosine-weighted sampling of the hemisphere?
- What is the Monte Carlo method (in general)?
- Explain how a Monte Carlo method can be used to solve the Rendering Equation.
- What is Russian Roulette and why do we need it?
- How can we use Russian Roulette and still have an unbiased estimator?

Lecture 20: Variance Reduction (Part 1 of 2)

- Be familiar with the following expression for Monte Carlo integration. What is the role of each term?

$$I = \lim_{n \to \infty} V(\Omega) \frac{1}{n} \sum_{i=1}^{n} f(X_i)$$

- Give an example of how we can reduce variance in our rendered results in a path tracing algorithm without increasing the number of samples.
- What does it mean for an estimator to be consistent?
- What does it mean for an estimator to be unbiased?
- Give a concrete example of how a renderer could give a biased estimate of an image. Is the renderer in your example consistent? Explain your answer.
- Give five examples of how you can reweight samples in a pathtracing algorithm in order to do importance sampling.
- What are the main ideas behind bidirectional path tracing?
- How would you enumerate all possible paths in a scene?
- How does Metropolis-Hastings sampling work?
- Assume you have code to generate random paths and code to mutate existing paths. Write pseudocode for Metropolis-Hastings path tracing.

Lecture 20: Variance Reduction (Part 2 of 2)

- Is this algorithm consistent? Is it unbiased? Is it efficient? For what kinds of scenes would this algorithm be best suited? Explain your reasoning for your answers to all of these questions.
- What is Stratified Sampling?
- Why is it preferred to random sampling?
- Hammersley and Halton points are pseudo random sampling techniques to generate points with low discrepancy. What is discrepancy? Why do we want to generate low discrepancy samples?
- Give a concise one sentence description of each of the following rendering algorithms that makes it clear the differences between them. Use a diagram to illustrate your description:
  - Rasterization
  - Ray casting
  - Ray tracing
  - Path tracing
  - Bidirectional path tracing
  - Metropolis Light Transport
  - Photon Mapping
  - Radiosity
- Which of these algorithms are best for capturing reflective surfaces? caustics? color bleeding? subsurface scattering? refraction? . . .

Lecture 21: Intro to Animation (Part 1 of 2)

- How were the first animations created? How were the first films created? Give some examples.
- Describe some of the first computer generated animations, giving the developer / artist and timeframe.
- Animations are created from keyframes. How do we interpolate between those keyframes?
- Why do we avoid splines of degree higher than three in computer graphics?
- Write a cubic polynomial P(t) in parameter t, which may describe a cubic spline.
- What are the constraints for P(t) to interpolate endpoints p1 at time t=0 and p2 at time t=1?
- What are the constraints for P(t) to have tangent vector r1 at time t=0 and r2 at time t=1?

Lecture 21: Intro to Animation (Part 2 of 2)

- This pair of constraints describes a Hermite spline. Derive the polynomial coefficients for the Hermite spline and write the cubic polynomial in terms of p1, p2, r1, and r2.
- Put your result in matrix form.
- Give properties of the Hermite spline in terms of continuity (C1, C2, etc..), interpolation, and local control.
- What type of spline has C2 continuity and interpolation, but not local control?
- What type of spline has C2 continuity and local control, but does not interpolate its key points?
- What are Catmull-Rom splines? How are tangents computed for Catmull-Rom splines?
- What are blend shapes and where are they used? What exactly is interpolated when using blend shapes for animation?

Lecture 22: Dynamics, Integration, and Optimization (Part 1 of 2)

- What is the "animation equation"?
- What is the difference between an ODE and a PDE? Give some examples of systems we can simulate by integrating an ODE.
- Sketch an overall system for simulating an ODE using a block diagram. Be clear about what is the state, how you advance the state forward in time, and what integrator you are choosing.
- When is the Euler-Lagrange equation useful?
- Be able to work through a simple example of obtaining dynamic equations of motion using Lagrangian mechanics.
- Describe how to put together a mass-spring system to simulate cloth.
- What are the forces on each cloth "particle"?
- What is Forward Euler integration and what is its disadvantage? Can you show a simple example where it fails?
- What is Backward Euler integration and what are its pros and cons?
- What is Symplectic Euler integration?

Lecture 22: Dynamics, Integration, and Optimization (Part 1 of 2)

- **Describe some problems in Computer Graphics where optimization is important.**
- **Be able to describe an optimization problem in standard form and give a couple of simple examples.**
- **How do you know you have a minimum of an objective function in an optimization problem without constraints? What properties must be true?**
- **What is meant by convexity of the domain in an optimization problem? convexity of the objective? Give examples of each. Why do we care about convexity in optimization?**

Lecture 23: Project 4 Overview

- Bezier, Hermite, and Catmull-Rom splines are really all the same thing. Any one representation can be converted to any of the others. Explain the differences between them.
- Be able to express a Hermite spline in different ways — as a cubic polynomial, in matrix form, or derive it from its control parameters.
- How do we ensure continuity between cubic spline segments? C0 continuity? .. C1 continuity? .. is C2 continuity possible in general?
- What is the difference between forward and inverse kinematics?
- Write an expression for the forward kinematics of a simple character or robot.
- What is the Jacobian? Compute the Jacobian for a simple character or robot.
- What is the Jacobian Transpose technique for inverse kinematics? Is it guaranteed to converge? What does that mean in practice? Does it give a locally optimal solution or a globally optimal one? Why?
- Be able to calculate positions for vertices on a mesh using linear blend skinning.
- Be able to use the grid version of the Laplacian to do smoothing on a grid.
- Be able to step the wave equation forward one time step using Euler integration.

Lecture 24: Optimization (Part 2) and PDEs

- **What is the difference between a PDE and an ODE? Give examples of when you might use each one and why.**
- **Interpret this sentence using an equation and a diagram: Solving a PDE looks like *"use neighbor information to get velocity (...and then add a little velocity each time)"***
- **Burger's equation is first order in time and second order in space. What does that mean? What are the orders of the Laplace equation? The heat equation? The wave equation? Be able to figure out the order of an equation from an expression of the equation itself.**
- **What are examples of questions we can answer using the Laplace equation, the heat equation, and the wave equation respectively?**
- **Outline the basic strategy for solving a PDE.**
- **What is the Laplace operator? Write it out as a sum of partial derivatives.**
- **Copy down the heat equation from the slides. Write out the process of solving this equation using Forward Euler integration.**
- **Copy down the wave equation from the slides. Write out the process of solving this equation using Forward Euler integration.**